

(12) **United States Patent**
Abbott, III et al.

(10) **Patent No.:** **US 9,443,037 B2**
(45) **Date of Patent:** ***Sep. 13, 2016**

(54) **STORING AND RECALLING INFORMATION TO AUGMENT HUMAN MEMORIES**

FOREIGN PATENT DOCUMENTS

(75) Inventors: **Kenneth H. Abbott, III**, Kirkland, WA (US); **Dan Newell**, Seattle, WA (US); **James O. Robarts**, Redmond, WA (US)

EP 0661627 A1 7/1995
EP 0759591 A1 2/1997

(Continued)

(73) Assignee: **Microsoft Technology Licensing, LLC**, Redmond, WA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 217 days.

Emma Berry; Using SenseCam, a wearable camera, to alleviate autobiographical memory loss; 2007; Microsoft Research Cambridge; pp. 38.*

(Continued)

This patent is subject to a terminal disclaimer.

Primary Examiner — Mariela Reyes

Assistant Examiner — Jermaine Mincey

(21) Appl. No.: **11/490,540**

(22) Filed: **Jul. 19, 2006**

(74) *Attorney, Agent, or Firm* — Sandy Swain; Micky Minhas

(65) **Prior Publication Data**

US 2007/0043459 A1 Feb. 22, 2007

Related U.S. Application Data

(63) Continuation of application No. 10/371,436, filed on Feb. 21, 2003, now Pat. No. 7,155,456, which is a continuation of application No. 09/876,814, filed on Jun. 6, 2001, now Pat. No. 6,549,915, which is a continuation of application No. 09/464,659, filed on Dec. 15, 1999, now Pat. No. 6,513,046.

(51) **Int. Cl.**

G06F 7/00 (2006.01)

G06F 17/30 (2006.01)

(52) **U.S. Cl.**

CPC ... **G06F 17/30997** (2013.01); **G06F 17/30017** (2013.01); **Y10S 707/99945** (2013.01); **Y10S 707/99948** (2013.01)

(58) **Field of Classification Search**

USPC 707/104.1
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

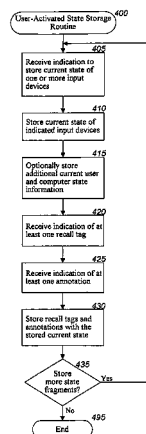
3,973,251 A 8/1976 Stephens
4,009,943 A 3/1977 Horimoto

(Continued)

(57) **ABSTRACT**

A system for computer-based storing of information about a current state so that later recall of the information can augment human memories. In particular, when information about a current event of interest is to be stored, a variety of current state information of different types (e.g., video, audio, and textual information) about the environment and about a user can be acquired via sensors and other input devices. The variety of state information can then be associated together as a group and stored for later retrieval. Other information can also be associated with the group, such as one or more recall tags that facilitate later retrieval of the group, or one or more annotations to provide contextual information when the other state information is later retrieved and presented to the user. When information about a past event is to be recalled, one or more identifying recall tags can be received that are used to identify one or more state information groups that match the identifying tags. Some or all of the previously-acquired state information for the identified state information groups can then be presented to the user on appropriate output devices. Other information, such as annotations, can also be presented to the user in order to describe the state information and thus assist the user's recollection of the previous state when the information was stored.

20 Claims, 19 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

4,283,712	A	8/1981	Goody	5,751,260	A	5/1998	Nappi et al.
4,458,331	A	7/1984	Amezcuca et al.	5,752,019	A *	5/1998	Rigoutsos et al.
4,569,026	A	2/1986	Best	5,754,938	A	5/1998	Herz et al.
4,815,030	A	3/1989	Cross et al.	5,761,662	A	6/1998	Dasan
4,864,501	A *	9/1989	Kucera et al. 704/8	5,769,085	A	6/1998	Kawakami et al.
4,874,963	A *	10/1989	Alspector 706/34	5,781,913	A	7/1998	Felsenstein et al. 707/501
4,901,096	A	2/1990	Lemelson	5,787,234	A	7/1998	Molloy
4,905,163	A	2/1990	Garber et al.	5,787,279	A *	7/1998	Rigoutsos 707/700
4,916,441	A	4/1990	Gombrich	5,790,974	A	8/1998	Tognazzini
4,970,683	A	11/1990	Harshaw et al.	5,794,235	A *	8/1998	Chess
4,991,087	A *	2/1991	Burkowski et al.	5,796,952	A	8/1998	Davis et al.
5,032,083	A	7/1991	Friedman 434/112	5,798,733	A	8/1998	Ethridge 342/357
5,133,075	A	7/1992	Risch	5,806,079	A *	9/1998	Rivette et al. 715/210
5,201,034	A	4/1993	Matsuura et al. 395/155	5,812,865	A	9/1998	Theimer et al. 395/800
5,208,449	A	5/1993	Eastman et al. 235/462	5,818,446	A	10/1998	Bertram et al.
5,214,757	A	5/1993	Mauney et al. 395/161	5,826,253	A	10/1998	Bredenberg
5,227,614	A	7/1993	Danielson et al. 235/380	5,831,594	A	11/1998	Tognazzini et al.
5,237,684	A	8/1993	Record et al.	5,832,296	A	11/1998	Wang et al.
5,251,294	A	10/1993	Abelow	5,835,087	A	11/1998	Herz et al.
5,267,147	A	11/1993	Harshaw et al.	5,852,814	A	12/1998	Allen
5,278,946	A	1/1994	Shimada et al.	5,867,171	A *	2/1999	Murata et al. 345/475
5,285,398	A	2/1994	Janik	5,873,070	A	2/1999	Bunte et al. 705/28
5,317,568	A	5/1994	Bixby et al.	5,878,274	A	3/1999	Kono et al. 395/828
5,327,529	A	7/1994	Fults et al.	5,879,163	A	3/1999	Brown et al.
5,335,276	A	8/1994	Thompson et al. 380/21	5,881,231	A *	3/1999	Takagi et al. 709/212
5,337,413	A *	8/1994	Lui et al. 710/2	5,899,963	A	5/1999	Hutchings
5,339,395	A	8/1994	Pickett et al.	5,902,347	A	5/1999	Backman et al. 701/200
5,349,654	A	9/1994	Bond et al.	5,905,492	A	5/1999	Straub et al.
5,353,399	A	10/1994	Kuwamoto et al.	5,910,765	A *	6/1999	Slemon et al. 340/517
5,388,198	A	2/1995	Layman et al.	5,910,799	A	6/1999	Carpenter et al. 345/333
5,398,021	A	3/1995	Moore	5,911,132	A	6/1999	Sloane
5,416,730	A	5/1995	Lookofsky 364/708.1	5,913,030	A	6/1999	Lotspiech et al.
5,454,074	A	9/1995	Hartel et al.	5,920,852	A *	7/1999	Graupe 706/26
5,470,233	A	11/1995	Fruchterman et al. 434/112	5,924,074	A	7/1999	Evans
5,471,629	A	11/1995	Risch	5,930,501	A *	7/1999	Neil 713/400
5,481,265	A	1/1996	Russell	5,937,160	A	8/1999	Davis et al.
5,481,667	A	1/1996	Bieniek et al.	5,938,721	A	8/1999	Dussell et al. 701/211
5,493,692	A	2/1996	Theimer et al. 455/26.1	5,942,986	A	8/1999	Shabot et al.
5,506,580	A	4/1996	Whiting et al. 341/51	5,945,988	A	8/1999	Williams et al.
5,513,646	A	5/1996	Lehrman et al.	5,948,041	A	9/1999	Abo et al.
5,522,024	A	5/1996	Hiraga et al.	5,953,718	A	9/1999	Wical
5,522,026	A	5/1996	Records et al.	5,959,611	A	9/1999	Smailagic et al. 345/156
5,535,323	A	7/1996	Miller et al.	5,963,914	A	10/1999	Skinner et al.
5,537,618	A	7/1996	Boulton et al.	5,963,940	A *	10/1999	Liddy et al.
5,539,665	A	7/1996	Lamming et al.	5,966,126	A	10/1999	Szabo
5,544,321	A	8/1996	Theimer et al.	5,966,533	A	10/1999	Moody
5,553,609	A	9/1996	Chen et al.	5,966,710	A *	10/1999	Burrows
5,555,376	A	9/1996	Theimer et al. 395/200.09	5,971,580	A	10/1999	Hall et al.
5,559,520	A	9/1996	Barzegar et al.	5,974,262	A	10/1999	Fuller et al.
5,560,012	A	9/1996	Ryu et al.	5,977,968	A	11/1999	LeBlanc
5,566,337	A	10/1996	Szymanski et al.	5,978,603	A	11/1999	Shiozaki et al.
5,568,645	A	10/1996	Morris et al. 395/800	5,980,096	A *	11/1999	Thalhammer-Reyero
5,572,401	A	11/1996	Carroll	5,983,335	A	11/1999	Dwyer, III 712/23
5,592,664	A	1/1997	Starkey	5,991,687	A	11/1999	Hale et al.
5,601,435	A	2/1997	Quy	5,991,735	A	11/1999	Gerace
5,603,054	A	2/1997	Theimer et al.	5,995,956	A	11/1999	Nguyen
5,611,050	A	3/1997	Theimer et al. 395/200.09	5,999,932	A	12/1999	Paul
5,642,303	A	6/1997	Small et al. 364/705.05	5,999,943	A *	12/1999	Nori et al.
5,646,629	A	7/1997	Loomis et al. 342/357	5,999,975	A	12/1999	Kittaka et al.
5,659,746	A *	8/1997	Bankert et al. 707/741	6,002,982	A	12/1999	Fry
5,675,358	A *	10/1997	Bullock et al. 345/420	6,002,994	A *	12/1999	Lane et al. 702/188
5,689,619	A	11/1997	Smyth	6,003,082	A	12/1999	Gampper et al.
5,689,708	A	11/1997	Regnier et al.	6,006,251	A	12/1999	Toyouchi et al.
5,701,894	A	12/1997	Cherry et al.	6,012,152	A	1/2000	Douik et al.
5,704,366	A	1/1998	Tacklind et al.	6,014,638	A	1/2000	Burge et al. 705/27
5,710,884	A	1/1998	Dedrick	6,023,729	A	2/2000	Samuel et al.
5,715,451	A	2/1998	Marlin	6,031,455	A	2/2000	Grube et al.
5,717,747	A	2/1998	Boyle et al.	6,035,264	A	3/2000	Donaldson et al.
5,719,744	A	2/1998	Jenkins et al. 361/683	6,041,331	A	3/2000	Weiner et al.
5,726,660	A	3/1998	Purdy et al. 342/357	6,041,365	A	3/2000	Kleinerman
5,726,688	A	3/1998	Siefert et al.	6,041,720	A	3/2000	Hardy
5,738,102	A	4/1998	Lemelson	6,044,415	A	3/2000	Futral et al.
5,740,037	A	4/1998	McCann et al.	6,047,301	A	4/2000	Bjorklund et al.
5,742,279	A *	4/1998	Yamamoto et al. 345/173	6,047,327	A	4/2000	Tso et al.
5,745,110	A	4/1998	Ertemalp	6,055,516	A	4/2000	Johnson et al.
				6,055,536	A	4/2000	Shimakawa et al.
				6,061,610	A	5/2000	Boer
				6,061,660	A	5/2000	Eggleston et al.
				6,064,943	A	5/2000	Clark, Jr. et al. 702/2

(56)

References Cited

U.S. PATENT DOCUMENTS

6,067,084 A	5/2000	Fado et al.	6,430,531 B1	8/2002	Polish
6,076,166 A	6/2000	Moshfeghi et al.	6,438,618 B1	8/2002	Lortz et al.
6,081,814 A	6/2000	Mangat et al.	6,442,549 B1	8/2002	Schneider
6,085,086 A	7/2000	La Porta et al.	6,442,589 B1	8/2002	Takahashi et al.
6,088,689 A	7/2000	Kohn et al.	6,442,620 B1	8/2002	Thatte et al.
6,091,411 A	7/2000	Straub et al.	6,446,076 B1	9/2002	Burkey et al.
6,092,101 A	7/2000	Birrell et al.	6,446,109 B2	9/2002	Gupta
6,094,625 A	7/2000	Ralston	6,460,036 B1	10/2002	Herz
6,098,065 A	8/2000	Skillen et al.	6,462,759 B1	10/2002	Kurtzberg et al.
6,105,063 A	8/2000	Hayes, Jr.	6,466,232 B1	10/2002	Newell et al.
6,108,197 A	8/2000	Janik 361/683	6,477,117 B1	11/2002	Narayanaswami et al.
6,108,665 A	8/2000	Bair et al.	6,483,485 B1	11/2002	Huang et al.
6,112,246 A	8/2000	Horbal et al.	6,484,200 B1	11/2002	Angal et al.
6,122,348 A	9/2000	French-St. George et al.	6,487,552 B1	11/2002	Lei et al.
6,122,657 A	9/2000	Hoffman, Jr. et al.	6,490,579 B1	12/2002	Gao et al.
6,122,960 A	9/2000	Hutchings et al.	6,505,196 B2	1/2003	Drucker et al.
6,127,990 A	10/2000	Zwern	6,507,567 B1	1/2003	Willars
6,128,663 A	10/2000	Thomas	6,507,845 B1	1/2003	Cohen et al.
6,131,067 A	10/2000	Girerd et al.	6,513,046 B1	1/2003	Abbott, III et al. 707/104.1
6,134,532 A	10/2000	Lazarus et al.	6,519,552 B1	2/2003	Sampath et al.
6,148,067 A *	11/2000	Leipow 379/201.01	6,526,035 B1	2/2003	Atarius et al.
6,154,745 A	11/2000	Kari et al.	6,529,723 B1	3/2003	Bentley
6,155,960 A	12/2000	Roberts et al.	6,539,336 B1	3/2003	Vock et al.
6,157,982 A	12/2000	Deo et al.	6,542,889 B1	4/2003	Aggarwal et al.
6,164,541 A	12/2000	Dougherty et al.	6,546,005 B1	4/2003	Berkley et al.
6,169,976 B1	1/2001	Colosso	6,546,425 B1	4/2003	Hanson et al.
6,185,534 B1	2/2001	Breese et al.	6,546,554 B1	4/2003	Schmidt et al.
6,188,399 B1	2/2001	Voas et al.	6,549,915 B2	4/2003	Abbott, III et al. 707/104.1
6,195,622 B1	2/2001	Altschuler et al.	6,549,944 B1	4/2003	Weinberg et al.
6,198,394 B1	3/2001	Jacobsen et al.	6,553,336 B1	4/2003	Johnson et al.
6,199,099 B1	3/2001	Gershman et al.	6,558,050 B1	5/2003	Ishibashi
6,199,102 B1	3/2001	Cobb	6,563,430 B1	5/2003	Kemink et al.
6,215,405 B1	4/2001	Handley et al.	6,568,595 B1	5/2003	Wiltz, Sr. et al.
6,218,958 B1	4/2001	Eichstaedt et al.	6,571,279 B1	5/2003	Herz et al.
6,223,190 B1 *	4/2001	Aihara et al. 715/234	6,578,019 B1	6/2003	Suda et al.
6,230,111 B1	5/2001	Mizokawa	6,615,197 B1	9/2003	Chai
6,236,768 B1	5/2001	Rhodes et al.	6,625,135 B1	9/2003	Johnson et al.
6,256,633 B1	7/2001	Dharap	6,629,133 B1	9/2003	Philyaw et al.
6,262,720 B1	7/2001	Jeffrey et al.	6,631,263 B1	10/2003	Corkery
6,263,268 B1	7/2001	Nathanson	6,636,831 B1	10/2003	Profit et al.
6,263,317 B1	7/2001	Sharp et al.	6,643,682 B1	11/2003	Todd et al.
6,272,470 B1	8/2001	Teshima	6,643,684 B1	11/2003	Malkin et al.
6,272,507 B1	8/2001	Pirolli et al.	6,652,283 B1	11/2003	Van Schaack et al.
6,282,517 B1	8/2001	Wolfe et al.	6,661,437 B1	12/2003	Miller et al.
6,282,537 B1	8/2001	Madnick et al.	6,672,506 B2	1/2004	Swartz et al.
6,285,757 B1	9/2001	Carroll et al.	6,697,836 B1	2/2004	Kawano et al.
6,285,889 B1	9/2001	Nykanen et al.	6,704,722 B2	3/2004	Wang Baldonado
6,289,316 B1	9/2001	Aghili et al.	6,704,785 B1	3/2004	Koo et al.
6,289,513 B1	9/2001	Bentwich	6,704,812 B2	3/2004	Bakke et al.
6,292,796 B1	9/2001	Drucker et al.	6,707,476 B1	3/2004	Hochstedler
6,294,953 B1	9/2001	Steeves 329/341	6,712,615 B2	3/2004	Martin
6,301,609 B1	10/2001	Aravamudan et al.	6,714,977 B1	3/2004	Fowler et al.
6,305,007 B1	10/2001	Mintz	6,718,332 B1	4/2004	Sitaraman et al.
6,305,221 B1	10/2001	Hutchings	6,732,273 B1	5/2004	Byers
6,308,203 B1	10/2001	Itabashi et al.	6,738,040 B2	5/2004	Jahn et al.
6,311,162 B1	10/2001	Reichwein et al.	6,738,759 B1	5/2004	Wheeler et al.
6,314,384 B1	11/2001	Goetz	6,738,975 B1	5/2004	Yee et al.
6,317,718 B1	11/2001	Fano	6,741,188 B1	5/2004	Miller et al.
6,321,158 B1	11/2001	DeLorme et al.	6,741,610 B1	5/2004	Voltsun et al.
6,321,279 B1	11/2001	Bonola	6,747,675 B1	6/2004	Abbott et al.
6,324,569 B1	11/2001	Ogilvie et al.	6,751,620 B2	6/2004	Orbanes et al.
6,327,535 B1	12/2001	Evans et al.	6,766,245 B2	7/2004	Padmanabhan
6,349,307 B1	2/2002	Chen	D494,584 S	8/2004	Schlieffers et al.
6,353,398 B1	3/2002	Amin et al.	6,791,580 B1	9/2004	Abbott et al.
6,353,823 B1	3/2002	Kumar	6,795,806 B1	9/2004	Lewis et al.
6,356,905 B1	3/2002	Gershman et al.	6,796,505 B2	9/2004	Pellaumail et al.
6,363,377 B1	3/2002	Kravets et al.	6,801,223 B1	10/2004	Abbott et al.
6,385,589 B1	5/2002	Trusheim et al.	6,812,937 B1	11/2004	Abbott et al.
6,392,670 B1	5/2002	Takeuchi et al.	6,829,639 B1	12/2004	Lawson et al.
6,401,085 B1	6/2002	Gershman et al.	6,834,195 B2	12/2004	Brandenberg et al.
6,405,159 B2	6/2002	Bushey et al.	6,834,208 B2	12/2004	Gonzales et al.
6,405,206 B1	6/2002	Kayahara	6,837,436 B2	1/2005	Swartz et al.
6,418,424 B1	7/2002	Hoffberg et al.	6,842,877 B2	1/2005	Robarts et al.
6,421,700 B1	7/2002	Holmes et al.	6,850,252 B1	2/2005	Hoffberg
6,427,142 B1	7/2002	Zachary et al.	6,853,966 B2	2/2005	Bushey et al.
			6,868,525 B1	3/2005	Szabo
			6,874,017 B1	3/2005	Inoue et al.
			6,874,127 B2	3/2005	Newell et al.
			6,885,734 B1	4/2005	Eberle et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

6,899,539 B1 5/2005 Stallman et al.
 6,963,899 B1 11/2005 Fernandez et al.
 6,968,333 B2 11/2005 Abbott et al.
 7,000,187 B2 2/2006 Messinger et al.
 7,010,501 B1 3/2006 Roslak et al.
 7,010,603 B2 3/2006 Martin, Jr. et al.
 7,040,541 B2 5/2006 Swartz et al.
 7,046,263 B1 5/2006 Abbott et al.
 7,055,101 B2 5/2006 Abbott et al.
 7,058,893 B2 6/2006 Abbott et al.
 7,058,894 B2 6/2006 Abbott et al.
 7,062,715 B2 6/2006 Abbott et al.
 7,063,263 B2 6/2006 Swartz et al.
 7,076,737 B2 7/2006 Abbott et al.
 7,080,322 B2 7/2006 Abbott et al.
 7,089,497 B2 8/2006 Abbott et al.
 7,096,253 B2 8/2006 Vinson et al.
 7,103,806 B1 9/2006 Horvitz
 7,107,516 B1* 9/2006 Anderson et al. 715/210
 7,107,539 B2 9/2006 Abbott et al.
 7,110,764 B1 9/2006 Blair et al.
 7,120,558 B2 10/2006 McIntyre et al.
 7,124,125 B2 10/2006 Cook et al.
 7,137,069 B2 11/2006 Abbott et al.
 7,143,093 B1 11/2006 Bracho et al.
 7,155,456 B2 12/2006 Abbott, III et al.
 7,162,473 B2 1/2007 Dumais et al.
 7,171,378 B2 1/2007 Petrovich et al.
 7,195,157 B2 3/2007 Swartz et al.
 7,203,906 B2 4/2007 Abbott et al.
 7,225,229 B1 5/2007 Abbott et al.
 7,231,439 B1 6/2007 Abbott et al.
 7,260,453 B2 8/2007 Poier et al.
 7,349,894 B2 3/2008 Barth et al.
 7,360,152 B2 4/2008 Capps et al.
 7,385,501 B2 6/2008 Miller et al.
 7,386,477 B2 6/2008 Fano
 7,392,486 B1 6/2008 Gyde et al.
 7,395,221 B2 7/2008 Doss et al.
 7,444,594 B2 10/2008 Abbott et al.
 7,464,153 B1 12/2008 Abbott et al.
 7,512,889 B2 3/2009 Newell et al.
 7,533,052 B2 5/2009 Tilfors et al.
 7,533,082 B2 5/2009 Abbott et al.
 7,561,200 B2 7/2009 Garvey, III et al.
 7,571,218 B2 8/2009 Tanaka et al.
 7,613,695 B1* 11/2009 Solomon et al.
 7,614,001 B2 11/2009 Abbott et al.
 7,647,400 B2 1/2010 Abbott et al.
 7,689,919 B2 3/2010 Abbott et al.
 7,734,780 B2 6/2010 Abbott et al.
 7,739,607 B2 6/2010 Abbott et al.
 7,779,015 B2 8/2010 Abbott et al.
 7,827,281 B2 11/2010 Abbott et al.
 7,877,686 B2 1/2011 Abbott et al.
 7,945,859 B2 5/2011 Abbott et al.
 8,020,104 B2 9/2011 Robarts et al.
 8,103,665 B2 1/2012 Abbott et al.
 8,626,712 B2 1/2014 Abbott et al.
 2001/0030664 A1 10/2001 Shulman et al.
 2001/0040590 A1 11/2001 Abbott et al.
 2001/0040591 A1 11/2001 Abbott et al.
 2001/0043231 A1 11/2001 Abbott et al.
 2001/0043232 A1 11/2001 Abbott et al.
 2002/0032689 A1 3/2002 Abbott, III et al.
 2002/0044152 A1 4/2002 Abbott, III et al.
 2002/0052930 A1 5/2002 Abbott et al.
 2002/0052963 A1 5/2002 Abbott et al.
 2002/0054130 A1 5/2002 Abbott, III et al.
 2002/0054174 A1 5/2002 Abbott et al.
 2002/0078204 A1 6/2002 Newell et al.
 2002/0080155 A1 6/2002 Abbott et al.
 2002/0080156 A1 6/2002 Abbott et al.
 2002/0083025 A1 6/2002 Robarts et al.
 2002/0083158 A1 6/2002 Abbott et al.

2002/0087525 A1 7/2002 Abbott et al.
 2002/0099817 A1 7/2002 Abbott et al.
 2002/0147880 A1 10/2002 WangBaldonado
 2002/0159770 A1 10/2002 Moultrie, Jr.
 2002/0191034 A1 12/2002 Sowizral et al.
 2003/0025798 A1 2/2003 Grosvenor et al.
 2003/0046401 A1 3/2003 Abbott et al.
 2003/0154476 A1 8/2003 Abbott, III et al. 725/37
 2003/0186201 A1 10/2003 Martin
 2003/0229900 A1 12/2003 Reisman
 2004/0088328 A1* 5/2004 Cook et al. 707/104.1
 2004/0133600 A1* 7/2004 Homer 707/104.1
 2004/0186854 A1* 9/2004 Choi 707/104.1
 2004/0201500 A1 10/2004 Miller et al.
 2004/0215663 A1* 10/2004 Liu et al. 707/104.1
 2004/0267700 A1 12/2004 Dumais et al.
 2004/0267812 A1* 12/2004 Harris et al. 707/104.1
 2005/0027704 A1 2/2005 Hammond et al.
 2005/0034078 A1 2/2005 Abbott et al.
 2005/0066282 A1 3/2005 Abbott et al.
 2005/0086243 A1 4/2005 Abbott et al.
 2005/0160113 A1* 7/2005 Sipusic et al. 707/104.1
 2005/0165843 A1* 7/2005 Capps et al. 707/104.1
 2005/0193017 A1* 9/2005 Kim 707/104.1
 2005/0203430 A1* 9/2005 Williams et al. 600/513
 2005/0266858 A1 12/2005 Miller et al.
 2005/0272442 A1 12/2005 Miller et al.
 2006/0004680 A1 1/2006 Robarts et al.
 2006/0019676 A1 1/2006 Miller et al.
 2006/0136393 A1 6/2006 Abbott et al.
 2006/0259494 A1 11/2006 Watson et al.
 2007/0022384 A1 1/2007 Abbott et al.
 2007/0089067 A1 4/2007 Abbott et al.
 2007/0130524 A1 6/2007 Abbott et al.
 2007/0168502 A1 7/2007 Abbott et al.
 2007/0185864 A1 8/2007 Budzik et al.
 2007/0266318 A1 11/2007 Abbott et al.
 2008/0090591 A1 4/2008 Miller et al.
 2008/0091537 A1 4/2008 Miller et al.
 2008/0147775 A1 6/2008 Abbott et al.
 2008/0161018 A1 7/2008 Miller et al.
 2008/0313271 A1 12/2008 Abbott et al.
 2009/0005079 A1 1/2009 Shields et al.
 2009/0013052 A1 1/2009 Robarts et al.
 2009/0055752 A1 2/2009 Abbott et al.
 2009/0094524 A1 4/2009 Abbott et al.
 2009/0150535 A1 6/2009 Abbott et al.
 2009/0228552 A1 9/2009 Abbott et al.
 2009/0234878 A1 9/2009 Herz et al.
 2009/0282030 A1 11/2009 Abbott et al.
 2010/0217862 A1 8/2010 Abbott et al.
 2010/0257235 A1 10/2010 Abbott et al.
 2010/0262573 A1 10/2010 Abbott et al.

FOREIGN PATENT DOCUMENTS

EP 0801342 A2 10/1997
 EP 0 823 813 A2 2/1998
 EP 0 846 440 A2 6/1998
 EP 0924615 A2 6/1999
 JP 05-260188 10/1993
 JP 09-091112 4/1997
 JP 11306002 11/1999
 WO WO 90/08361 7/1990
 WO WO95/31773 A1 11/1995
 WO WO 97/03434 1/1997
 WO WO97/31388 A2 9/1997
 WO WO9800787 1/1998
 WO WO 98/47084 10/1998
 WO WO 99/17228 4/1999
 WO WO99/26180 A1 5/1999
 WO WO99/66394 12/1999
 WO WO99/67698 A2 12/1999
 WO WO-00/36493 A1 6/2000

OTHER PUBLICATIONS

“‘Affective Understanding:’ Modeling and Responding to User Affect,” retrieved Oct. 2, 1998, from http://www.media.mit.edu/affect/AC_research/understanding.html, 3 pages.

(56)

References Cited

OTHER PUBLICATIONS

- "Alps GlidePoint," retrieved Oct. 2, 1998, from <http://www.alps.com/p17.html>, 1 page.
- "GyroPoint Technology," retrieved Oct. 2, 1998, from <http://www.gyration.com/html/gyropoint.html>, 3 pages.
- "Haptics," retrieved Oct. 2, 1998, from <http://www.ai.mit.edu/projects/handarm-haptics/haptics.html>, 2 pages.
- "Research Areas in Affective Computing," retrieved Oct. 2, 1998, from <http://www.media.mit.edu/affect/>, 1 page.
- "Research on Affective Pattern Recognition and Modeling," retrieved Oct. 2, 1998, from http://www.media.mit.edu/affect/AC_research/recognizing.html, 4 pages.
- "Research on Sensing Human Affect," retrieved Oct. 2, 1998, from http://www.media.mit.edu/affect/AC_research/sensing.html, 5 pages.
- "Smart Rooms," retrieved Oct. 2, 1998, from <http://vismod.www.media.mit.edu/vismod/demos/smartroom/>, 3 pages.
- "SmartDesk Home Page," retrieved Oct. 2, 1998, from <http://vismod.www.media.mit.edu/vismod/demos/smartdesk/>, 4 pages.
- "The MIT Wearable Computing Web Page," retrieved Oct. 2, 1998, from <http://wearables.www.media.mit.edu/projects/wearables/>, 3 pages.
- "Wearable Computer Systems for Affective Computing," retrieved Oct. 2, 1998, from http://www.media.mit.edu/affect/AC_research/wearables.html, 5 pages.
- Aoki, H., et al., "Realtime Personal Positioning System for a Wearable Computer," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 7 pages.
- Billingshurst, M., et al., "Wearable Devices: New Ways to Manage Information," *Computer*, 32(1):57-64, Jan. 1999, IEEE Computer Society, 8 pages.
- Bowskill, J., et al., "Wearable Location Mediated Telecommunications; A First Step Towards Contextual Communication," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 8 pages.
- Dey, A., et al., "The Conference Assistant: Combining Context-Awareness with Wearable Computing," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 7 pages.
- Golding, A., et al., "Indoor Navigation Using a Diverse Set of Cheap, Wearable Sensors," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 8 pages.
- Hull, R., et al., "Towards Situated Computing," Proceedings 1st International Symposium on Wearable Computers (ISWC'97), Cambridge, MA, Oct. 13-14, 1997, 8 pages.
- Kirsch, D., "The Sentic Mouse: A Tool for Measuring Emotional Valence," retrieved Oct. 2, 1998, from http://www.media.mit.edu/affect/AC_research/projects/sentic_mouse.html, 2 pages.
- Kortuem, G., et al., "When Cyborgs Meet: Building Communities of Cooperating Wearable Agents," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 9 pages.
- Lashkari, Y., et al., "Collaborative Interface Agents," Proceedings of the Twelfth National Conference on Artificial Intelligence (vol. 1), Seattle, Washington, Aug. 1994, 10 pages.
- Lehikoinen, J., et al., "MEX: A Distributed Software Architecture for Wearable Computers," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 6 pages.
- Leonhardi, A., et al., "Virtual Information Towers—A Metaphor for Intuitive, Location-Aware Information Access in a Mobile Environment," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 6 pages.
- Lunt, T., et al., "Knowledge-Based Intrusion Detection," Proceedings of the Annual Artificial Intelligence Systems in Government Conference Mar. 27-31, 1989, IEEE Comp. Soc. Press, 1989, pp. 102-107, 6 pages.
- Macs, P., "Agents that Reduce Work and Information Overload," *Communications of the ACM*, 37(7):30-40, Jul. 1994, 13 pages.
- Metz, C., "MIT: Wearable PC's, Electronic Ink, and Smart Rooms," *PC Magazine*, Jun. 9, 1998, pp. 192-193, 2 pages.
- Oakes, C., "The Truman Show Realized?," retrieved Oct. 21, 1998, from <http://www.wired.com/news/news/technology/story/15745.html>, 4 pages.
- Picard, R., et al., "Affective Wearables," Proceedings 1st International Symposium on Wearable Computers (ISWC'97), Cambridge, MA, Oct. 13-14, 1997, 8 pages.
- Rekimoto, J. et al., "The World Through the Computer: Computer Augmented Interaction with Real World Environments," Symposium on User Interface Software and Technology (IST '95) ACM Press, Nov. 1995, pp. 29-36, 14 pages.
- Rhodes, B., "The Wearable Remembrance Agent: A System for Augmented Memory," Proceedings 1st International Symposium on Wearable Computers (ISWC'97), Cambridge, MA, Oct. 13-14, 1997, 8 pages.
- Rhodes, B., "WIMP Interface Considered Fatal," retrieved Jul. 23, 1998, from <http://rhodes.www.media.mit.edu/people/rhodes/Papers/no-wimp.html>, 3 pages.
- Rogers, E. et al., "Outbreak Agent: Intelligent Wearable Technology for Hazardous Environments," 1997 IEEE International Conference on 'Computational Cybernetics and Simulation', Oct. 12-15, 1997, pp. 3198-3203, 8 pages.
- Sato, J., et al., "Autonomous Behavior Control of Virtual Actors Based on the AIR Model," Proceedings Computer Animation, Jun. 5, 1997, pp. 113-118, 5 pages.
- Schneider, J., et al., "Modeling Wearable Negotiation in an Opportunistic Task Oriented Domain," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 2 pages.
- Smailagic, A., et al., "MoCCA: A Mobile Communications and Computing Architecture," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 8 pages.
- Tan, H., et al., "Tactical Displays for Wearable Computing," Proceedings 1st International Symposium on Wearable Computers (ISWC'97), Cambridge, MA, Oct. 13-14, 1997, pp. 84-88, 5 pages.
- Yang, J., et al., "Smart Sight: A Tourist Assistant System," Proceedings 3rd International Symposium on Wearable Computers (ISWC'99), San Francisco, CA, Oct. 18-19, 1999, 6 pages.
- Amon, et al., "Integration of Design Education, Research and Practice at Carnegie Mellon University: A Multi-Disciplinary Course in Wearable Computer Design," Proceedings of The Frontiers in Education Conference, Nov. 1-4, 1995, pp. 4a1.14-4a1.22, vol. 2.
- Bacon, et al., "Using Events to Build Distributed Applications," University of Cambridge, 1996. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.34.4545>. Last accessed Dec. 9, 2008, 8 pages.
- Bier, et al., "Toolglass and Magic Lenses: The See-Through Interface," Proceedings of SIGGRAPH'93, Computer Graphics and Annual Conference Series, ACM, pp. 73-80, Anaheim, California, 1993. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.36.9741>. Last accessed Dec. 9, 2008, 8 pages.
- Billingshurst, et al., "An Evaluation of Wearable Information Spaces," Proceedings of the Virtual Reality Annual International Symposium, 1998, 8 pages.
- Billingshurst, Research Directions in Wearable Computing, University of Washington, May 1998, 48 pages.
- Bishop, "Hard Ware," Jul. 28, 2000, Puget Sound Business Journal, print edition, pp. 1-4.
- "BridgeVIEW and LabVIEW G Programming Reference Manual," Jan. 1998, National Instruments Corporation. <http://www.ni.com/pdf/manuals/321296b.pdf>. Last accessed Dec. 7, 2008, 667 pages.
- Brown, et al., "Using Explicit Requirement and Metrics for Interface Agent User Model Correction," Autonomous Agents '98.
- Brown, et al., "Utility Theory-Based User Models for Intelligent Interface Agents," Proceedings of the Twelfth Canadian Conference on Artificial Intelligence (AI'98), Jun. 1998.

(56)

References Cited

OTHER PUBLICATIONS

- Budzik, et al. "Watson: Anticipating and Contextualizing Information Needs," May 1, 1999, Proceedings of the 62nd Annual Meeting of the American Society for Information Science, pp. 1-14.
- T. Joachims, Text categorization with support vector machines: learning with many relevant features, Machine Learning, European Conference on Machine Learning, Apr. 21, 1998, pp. 137-142.
- Theimer, et al. Operating System Issues for PDAs, In Fourth Workshop on Workstation Operating Systems, 1993, 7 pages.
- Wachowicz, et al. "Combining Location and Data Management in an Environment for Total Mobility," University of Cambridge, England, 1996. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.37.4550>. Last accessed Dec. 9, 2008, 12 pages.
- Want, Active Badges and Personal Interactive Computing Objects, IEEE Transactions on Consumer Electronics, 1992, 11 pages, vol. 38—No. 1.
- Want, et al., The Active Badge Location System, ACM Transactions on Information Systems, Jan. 1992, pp. 91-102, vol. 10—No. 1.
- Wardell. "Stardock.net Releases WindowBlinds." Stardock.net, Inc. Oct. 4, 1999. http://stardock.net/media/pr_wb10.html. Last accessed Dec. 9, 2008, 2 pages.
- Weiser. The Computer for the 21st Century, Scientific American, Sep. 1991, 8 pages.
- Weiser. Some Computer Science Issues in Ubiquitous Computing, Communications of the ACM, Jul. 1993, pp. 75-84, vol. 36—No. 7.
- Zenel, et al. "A General Purpose Proxy Filtering Mechanism Applied to the Mobile Environment," Proceedings of the 3rd annual ACM/IEEE international conference on Mobile computing and networking, MOBICOM'97, Budapest, Hungary, pp. 248-259. <http://portal.acm.org/citation.cfm?id=262153>. Last accessed Dec. 9, 2008, 12 pages.
- Kortuem, et al. Context Aware, Adaptive, Wearable Computers as Remote Interfaces to Intelligent Environments. University of Oregon, Oct. 1998, 8 pages.
- "LabVIEW User Manual," Jan. 1998 Edition, National Instruments. <http://claymore.engineer.gvsu.edu/eod/courses/egr450/media/320999b.pdf>. Last accessed Dec. 7, 2008, 514 pages.
- Losee, Jr. Minimizing information overload: the ranking of electronic messages, Journal of Information Science 15, Elsevier Science Publishers B.V., 1989, pp. 179-189.
- Mann. Smart Clothing: Wearable Multimedia Computing and Personal Imaging to Restore the Technological Balance Between People and Their Environments. ACM Multimedia, Nov. 1996, pp. 163-174.
- Mott, et al. "A Formalism for Context Mediation Based on Feature Logic." Feb. 1998. <http://citeseem.ist.psu.edu/viewdoc/summary?doi=10.1.1.52.4473>. Last accessed Dec. 10, 2008, 11 pages.
- Ni. "Attribute Name Evaluation and Its Implementation," School of Computing and Mathematics, Deakin University, Geelong, Victoria, May 16, 1994. <http://www.deakin.edu.au/scitech/sit/dsapp/archive/techreport/TR-C94-10.pdf>. Last accessed Dec. 9, 2008, 32 pages.
- OA Dated Sep. 15, 2008 for U.S. Appl. No. 11/033,974, 17 pages.
- OA Dated Aug. 14, 2008 for U.S. Appl. No. 10/981,941, 14 pages.
- OA Dated Dec. 3, 2008 for U.S. Appl. No. 09/981,320, 40 pages.
- OA Dated Aug. 1, 2008 for U.S. Appl. No. 11/179,822, 17 pages.
- OA Dated Sep. 7, 2008 for U.S. Appl. No. 09/894,642, 28 pages.
- OA Dated Jul. 29, 2008 for U.S. Appl. No. 10/984,511, 28 pages.
- OA Dated Dec. 4, 2008 for U.S. Appl. No. 11/567,902, 21 pages.
- OA Dated Nov. 8, 2008 for U.S. Appl. No. 11/761,210, 11 pages.
- Ockerman, et al. "Wearable Computer for Performance Support: Initial Feasibility Study," International Symposium in Wearable Computers, Oct. 1997, pp. 10-17.
- Papakonstantinou, et al. "MedMaker: A Mediation System Based on Declarative Specifications." 1995. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.35.735>. Last accessed Nov. 25, 2008, 10 pages.
- Rhodes, Remembrance Agent: A continuously running automated information retrieval system, The Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi Agent Technology, 1996, pp. 487-495.
- Rhodes, The Wearable Remembrance Agent: A System for Augmented Memory, Personal Technologies Journal Special Issue on Wearable Computing, 1997, 12 pages.
- Schilit. A System Architecture for Context-Aware Mobile Computing, Columbia University, 1995, 153 pages.
- Schilit, et al., The ParcTab Mobile Computing System, IEEE WWOS-IV, 1993, 4 pages.
- Schilit, et al., Context-Aware Computing Applications, In Proceedings of the Workshop on Mobile Computing Systems and Applications, Dec. 1994, pp. 85-90.
- Schilit, et al., Customizing Mobile Applications, Proceedings USENIX Symposium on Mobile and Location Independent Computing, Aug. 1993, 9 pages.
- Schilit, et al., Disseminating Active Map Information to Mobile Hosts, IEEE Network, 1994, pp. 22-32, vol. 8—No. 5.
- Schmidt, et al. "There's More to Context Than Location." Proceedings of the International Workshop on Interactive Applications of Mobile Computing (IMC98), Germany, Nov. 1998, 10 pages. http://citeseer.comp.nus.edu.sg/cache/papers/cs/16114/http://zSzzSzwwww.teco.edu/Sz-albrechtSzpublicationzSzdraft_docszSzcontext-is-more-than-location.pdf/there-is-more-to.pdf.
- Smailagic, et al. MoCCA: A Mobile Communication and Computing Architecture. Third International Symposium on Wearable Computers, Oct. 18-19, 1999, San Francisco, California.
- Spreitzer, et al. "Providing Location Information in a Ubiquitous Computing Environment," Xerox Palo Alto Research Center, ACM, 1993, pp. 270-283. <http://www.comp.lancs.ac.uk/computing/staff/kc/Lecturing/MSc/wk10/p270-spreitzer.pdf>. Last accessed Dec. 9, 2008, 14 pages.
- Spreitzer, et al. Architectural Considerations for Scalable, Secure, Mobile Computing with Location Information, In The 14th International Conference on Distributed Computing Systems, Jun. 1994, pp. 29-38.
- Spreitzer et al. Scalable, Secure, Mobile Computing with Location Information, Communications of the ACM, Jul. 1993, 1 page, vol. 36—No. 7.
- Starner. Wearable Computing and Contextual Awareness, Massachusetts Institute of Technology, Jun. 1999, 248 pages.
- OA dated Jan. 2, 2009 for U.S. Appl. No. 11/559,854, 23 pages.
- Chen, et al., A Survey of Context-Aware Mobile Computing Research, Dartmouth Computer Science Technical Report, 2000, 16 pages.
- "Context-Awareness in Wearable and Ubiquitous Computing." Future Computing Environments, 1997. GVVU Technical Report GIT-GVVU-97-11. <http://www.cc.gatech.edu/fce/pubs/iswc97/wear.html>. Last accessed Dec. 8, 2008, 13 pages.
- "Context Recognition by User Situation Data Analysis (Context)." <http://www.cs.helsinki.fi/group/context/>. Last accessed Dec. 19, 2008, 7 pages.
- Crabtree, et al. "Wearable Computing and the Remembrance Agent," BT Technology Journal, vol. 16, No. 3, Jul. 1998, pp. 118-124. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.31.8514>. Last accessed Dec. 9, 2008, 7 pages.
- D'Orazio, et al. Mobile Robot Position Determination Using Visual Landmarks. IEEE Transactions on Industrial Electronics, vol. 41, issue 6, Dec. 1994, pp. 654-662. Last accessed Dec. 9, 2008, 9 pages.
- Doorenbos, et al. "A Scalable Comparison-Shopping Agent for the World-Wide-Web," Proceedings of the First International Conference on Autonomous Agents, Marina Del Rey, California, Feb. 5-8, 1997, pp. 39-48. <http://www.cs.washington.edu/homes/etzioni/papers/agents97.pdf>. Last accessed Dec. 9, 2008, 10 pages.
- Finger, et al. Rapid Design and Manufacture of Wearable Computers. Communication of the ACM, vol. 39, No. 2, Feb. 1996, pp. 63-68.
- Goh, et al. "Context Interchange: New Features and Formalisms for the Intelligent Integration of Information," ACM Transactions on Information Systems, 1997. <http://dspace.mit.edu/bitstream/handle/1721.1/2642/SWP-3941-36987967.pdf?sequence=1>. Last accessed Dec. 10, 2008, 25 pages.
- Han, et al. "DBMiner: A System for Mining Knowledge in Large Relational Databases," Proceedings 1996 International Conference

(56)

References Cited

OTHER PUBLICATIONS

on Data Mining and Knowledge Discovery, Portland, OR, Aug. 1996. <http://www.cs.ualberta.ca/~zaiane/postscript/kdd96.pdf>. Last accessed Dec. 9, 2008, 6 pages.

Harter, et al., "A Distributed Location System for the Active Office," IEEE Network, 1994, pp. 62-70.

Horvitz, et al., In Pursuit of Effective Handsfree Decision Support: Coupling Bayesian Inference, Speech Understanding, and User Models, 1995, 8 pages.

Horvitz, et al., Attention-Sensitive Alerting in Computing Systems, Microsoft Research, Aug. 1999.

"Intelligent User Interface Prompt Level," IBM Technical Disclosure Bulletin, IBM Corp., New York, vol. 35, No. 1A, Jun. 1992, pp. 25-26.

International Search Report dated Sep. 29, 2003 for PCT Application Serial No. 00/20685, 3 Pages.

International Search Report PCT/US01/32543, Oct. 27, 2003, 3 pages.

International Search Report for PCT Application No. PCT/US01/10394, mailed Jun. 13, 2002, 5 pages.

Jakobovits, "Integrating Autonomous Heterogeneous Information Sources," University of Washington, Technical Report, UV-CSE-971205, Jul. 15, 1997, 29 pages.

Office Action dated Jun. 8, 2009 in U.S. Appl. No. 10/790,602.

International Search Report, Application No. PCT/US01/10599, Nov. 28, 2002.

Rhodes, The Wearable Remembrance Agent: A System for Augmented Memory. Personal Technologies Journal Special Issue on Wearable Computing, 1997, 12 pages.

Smailagic, et al., "Matching Interface Design with User Task: Modalities of Interaction with CMU Wearable Computers," IEEE Personal Communications, pp. 14-25, Feb. 1996.

Starner, et al., "Visual Contextual Awareness in Wearable Computing," Media Lab, MIT, Oct. 1998.

European Office Action dated Feb. 13, 2009 for EP Application No. 01964695.9, 4 pages.

Bates et al., Middleware Support for Mobile Multimedia Applications, ICL System Journal, Nov. 1997, 20 pages.

Cochran, "The Development of Wearable Clothing," Dr. Dobb's, online edition, pp. 1-2.

Cugola et al., Exploiting an Event-Based Infrastructure to Develop Complex Distributed Systems, 20th International Conference on Software Engineering (ICSE'98) Apr. 1998, p. 261-270, 10 pages.

Dechamboux et al., Using a Distributed Shared Memory for Implementing Efficient Information Mediators, Proceedings of the International Conference and Exhibition on High-Performance Computing and Networking, Apr. 1996, 5 pages.

Dey et al., CyberDesk: A Framework for Providing Self-Integrating Context-Aware Services, Knowledge-Based Systems, 11:3-13, 1998, 10 pages.

Goharian et al., Enterprise Text Processing: A Sparse Matrix Approach, 2001, IEEE, pp. 71-75.

International Search Report, PCT/US01/10538, Aug. 14, 2002, 2 pages.

OA Dated Oct. 30, 2008 for U.S. Appl. No. 11/490,540, 37 pages.

Ockerbloom, Mediating Among Diverse Data Formats: Thesis Summary, PhD Thesis, Technical Report CMU-CS-98-102, Department of Computer Sciences, Carnegie Mellon University, Pittsburgh, PA, Jan. 14, 1998, 15 pages.

Pascoe, Adding Generic Contextual Capabilities to Wearable Computers, Proceedings of the 2nd International Symposium on Wearable Computers (ISWC '98), Pittsburgh, PA, Oct. 19-20, 1998, 8 pages.

Sawhney, Contextual awareness, messaging and communication in nomadic audio environments, Massachusetts Institute of Technology, Jun. 1998 pp. 80-105, 25 pages.

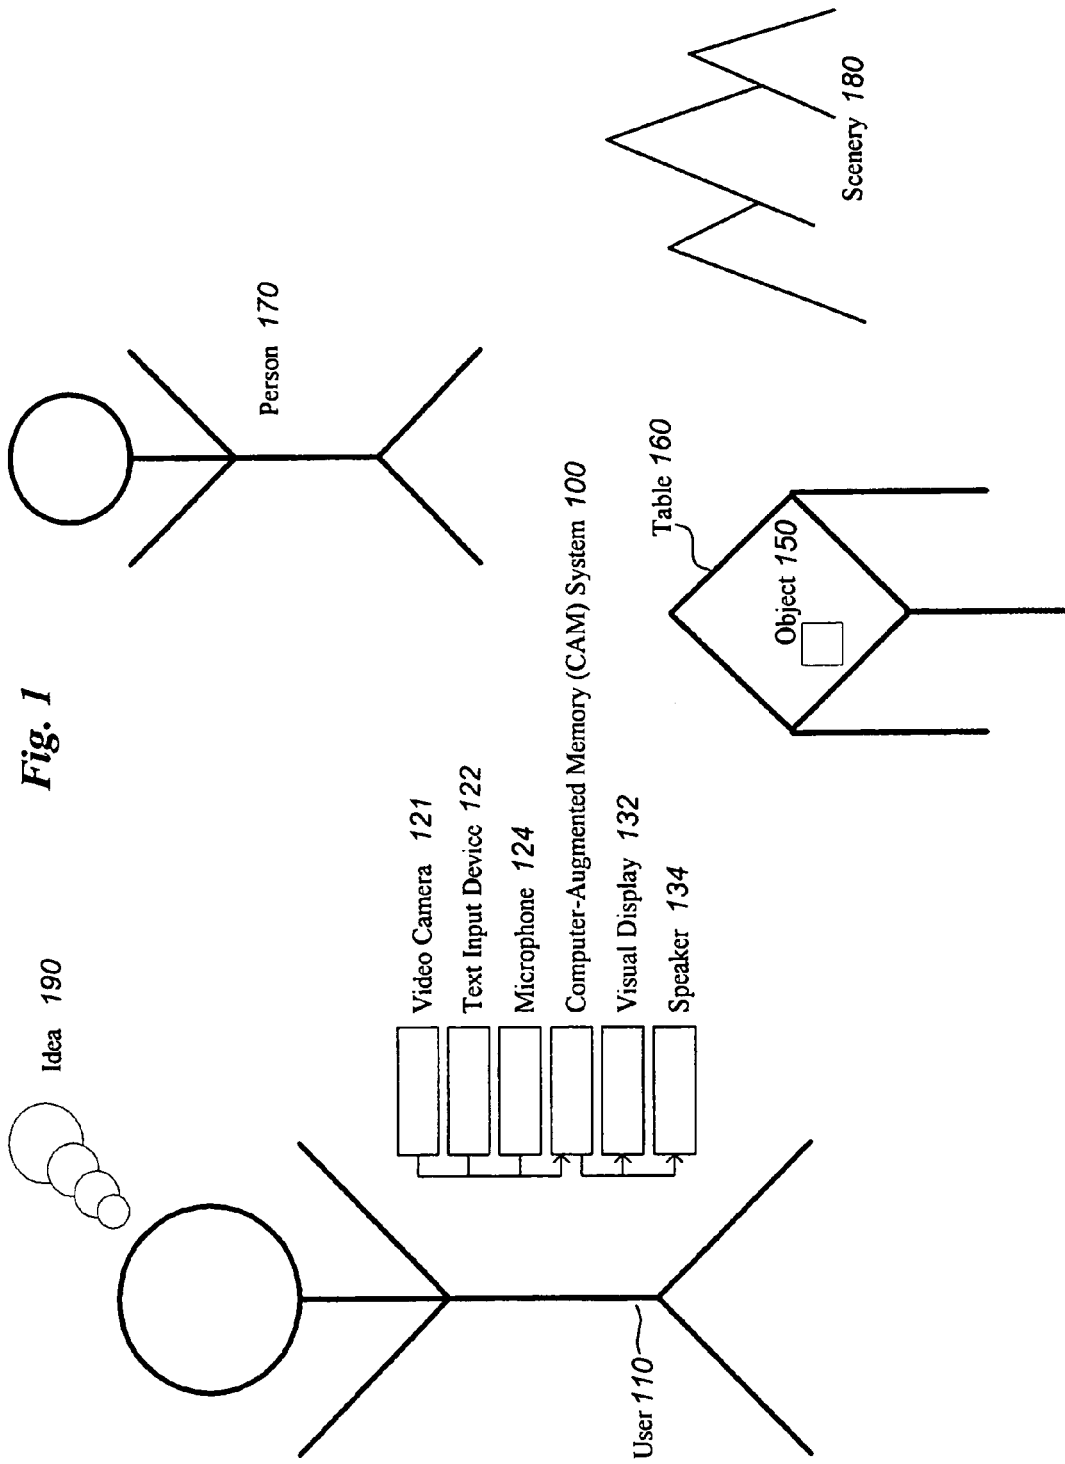
Starovic, An Event Based Object Model for Distributed Programming, Technical Report TCD-CS-95-28, Dept. of Computer Science, Trinity College Dublin, Ireland, 1995, retrieved from <http://citeseer.ist.psu.edu/starovic95event.html>, 15 pages.

Salber et al., The Design and Use of a Generic Context Server. Technical Report GIT-GVU-98-32. GVU Center, College of Computing, Georgia Institute of Technology. 1998. Available at <ftp://ftp.cc.gatech.edu/pub/gvu/tr/1998/98-32.pdf>. Last accessed Mar. 12, 2010.

European Search Report for EP 02017656.6 mailed Mar. 18, 2009.

Teranishi et al., "Dynamic Object Recomposition for Active Information System on Mobile Environment", In Proceedings of International Database Engineering and Applications Symposium, IEEE, 1997, pp. 220-228.

* cited by examiner



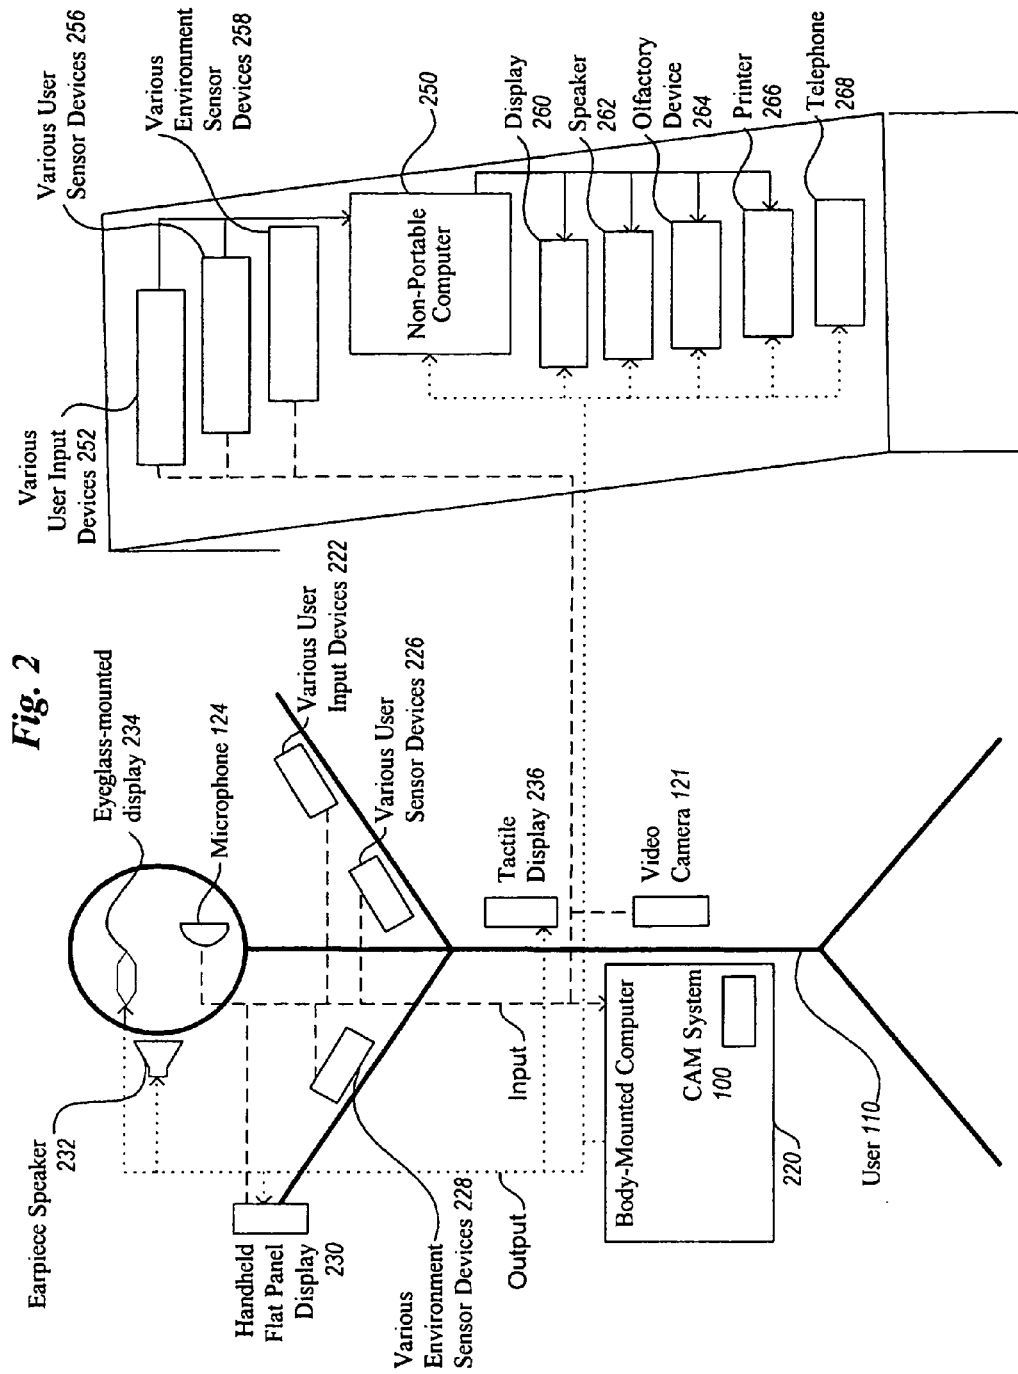
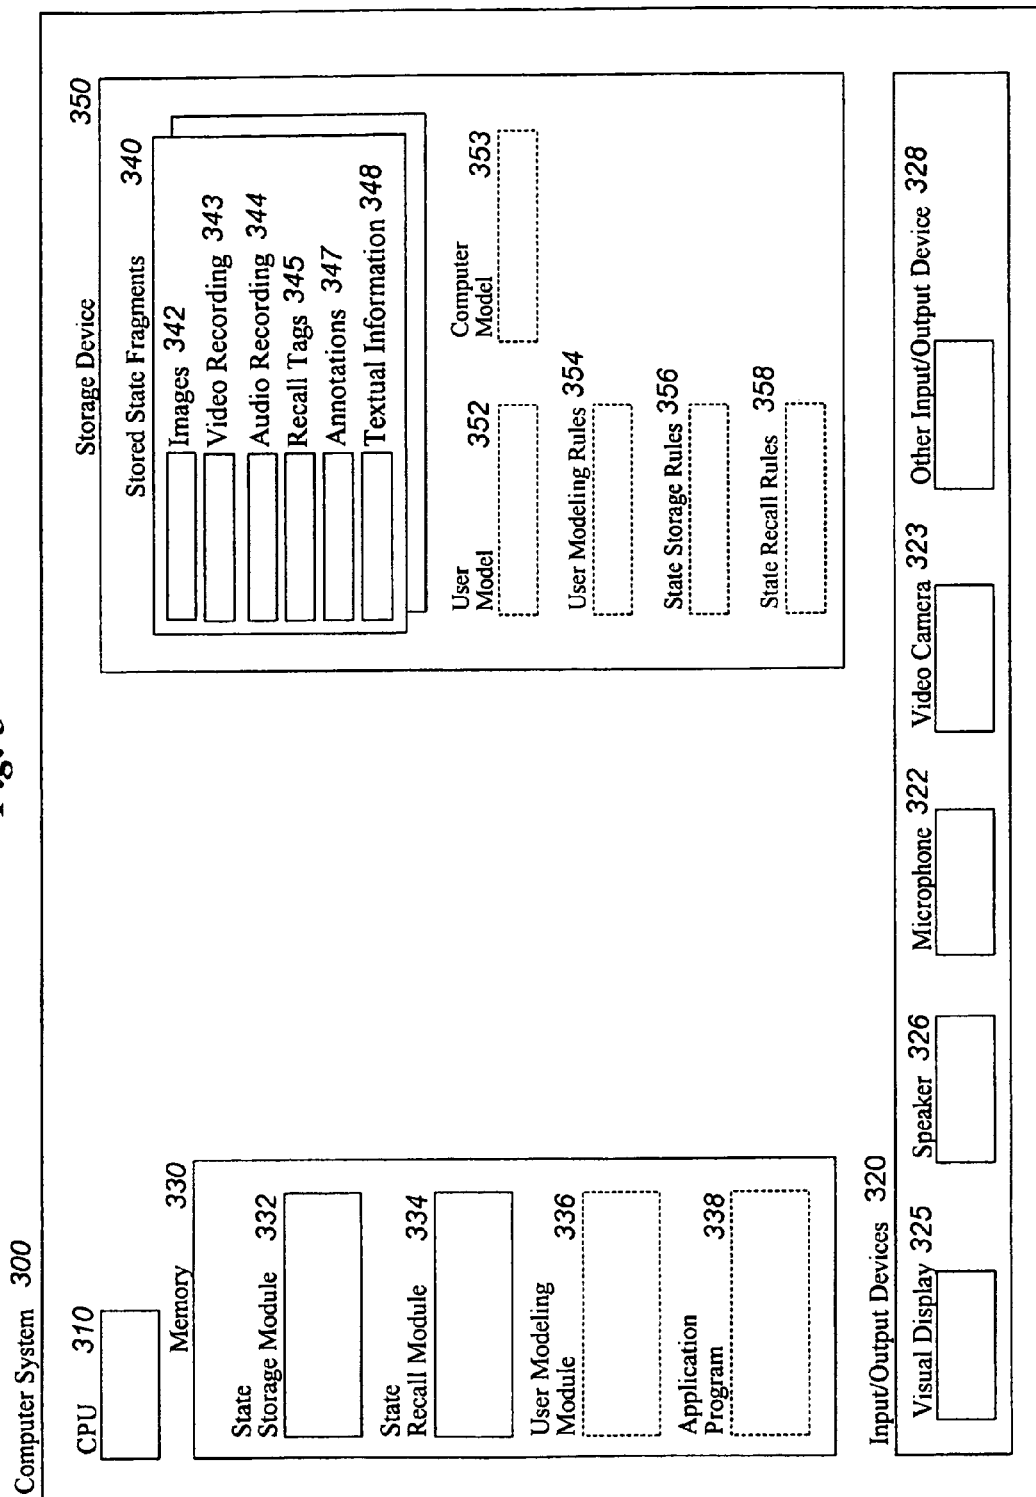


Fig. 3



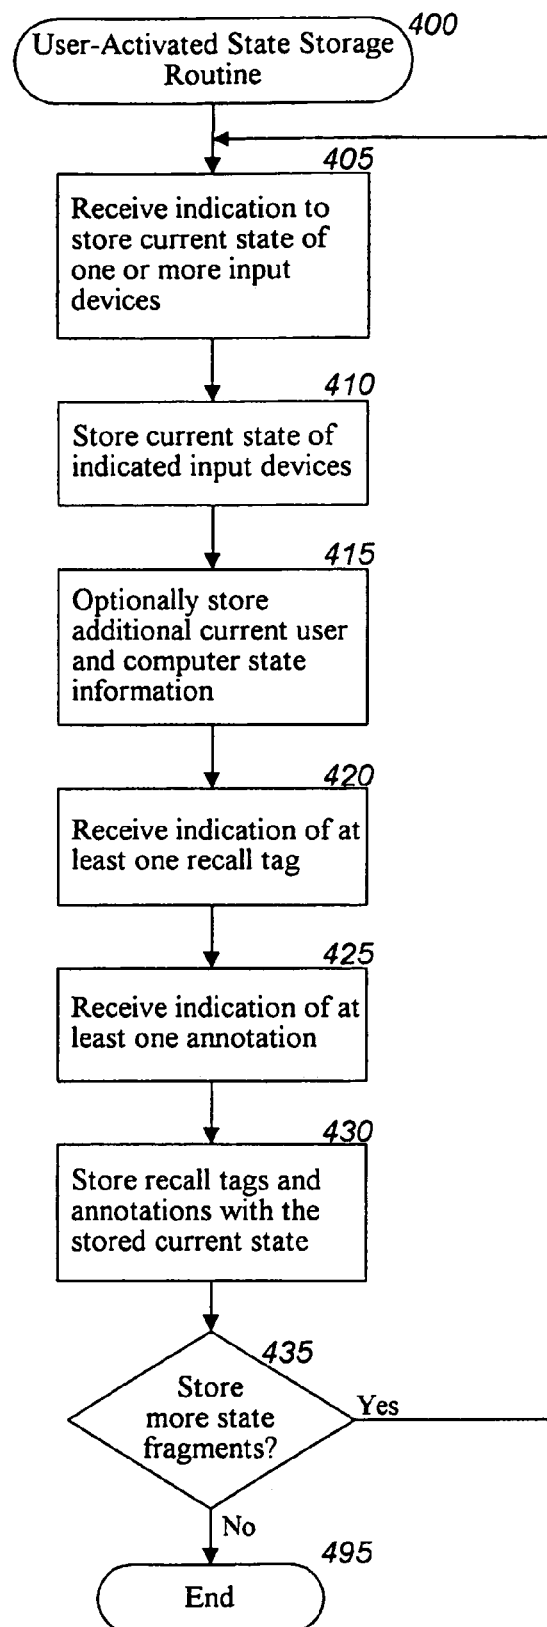
*Fig. 4*

Fig. 5

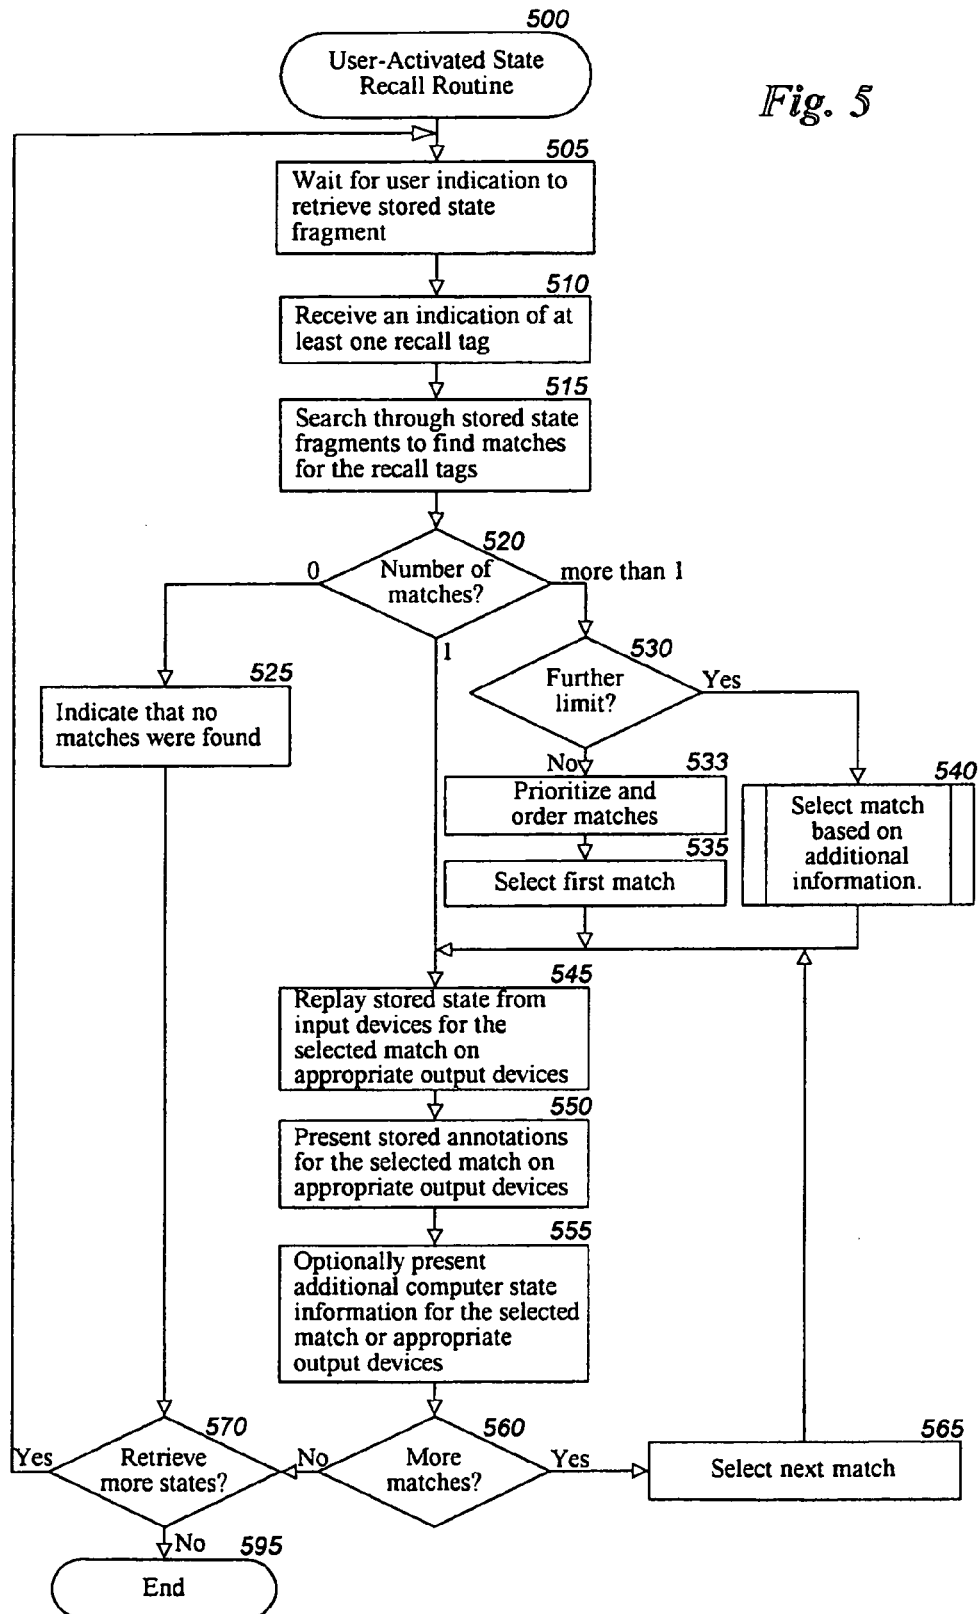


Fig. 6

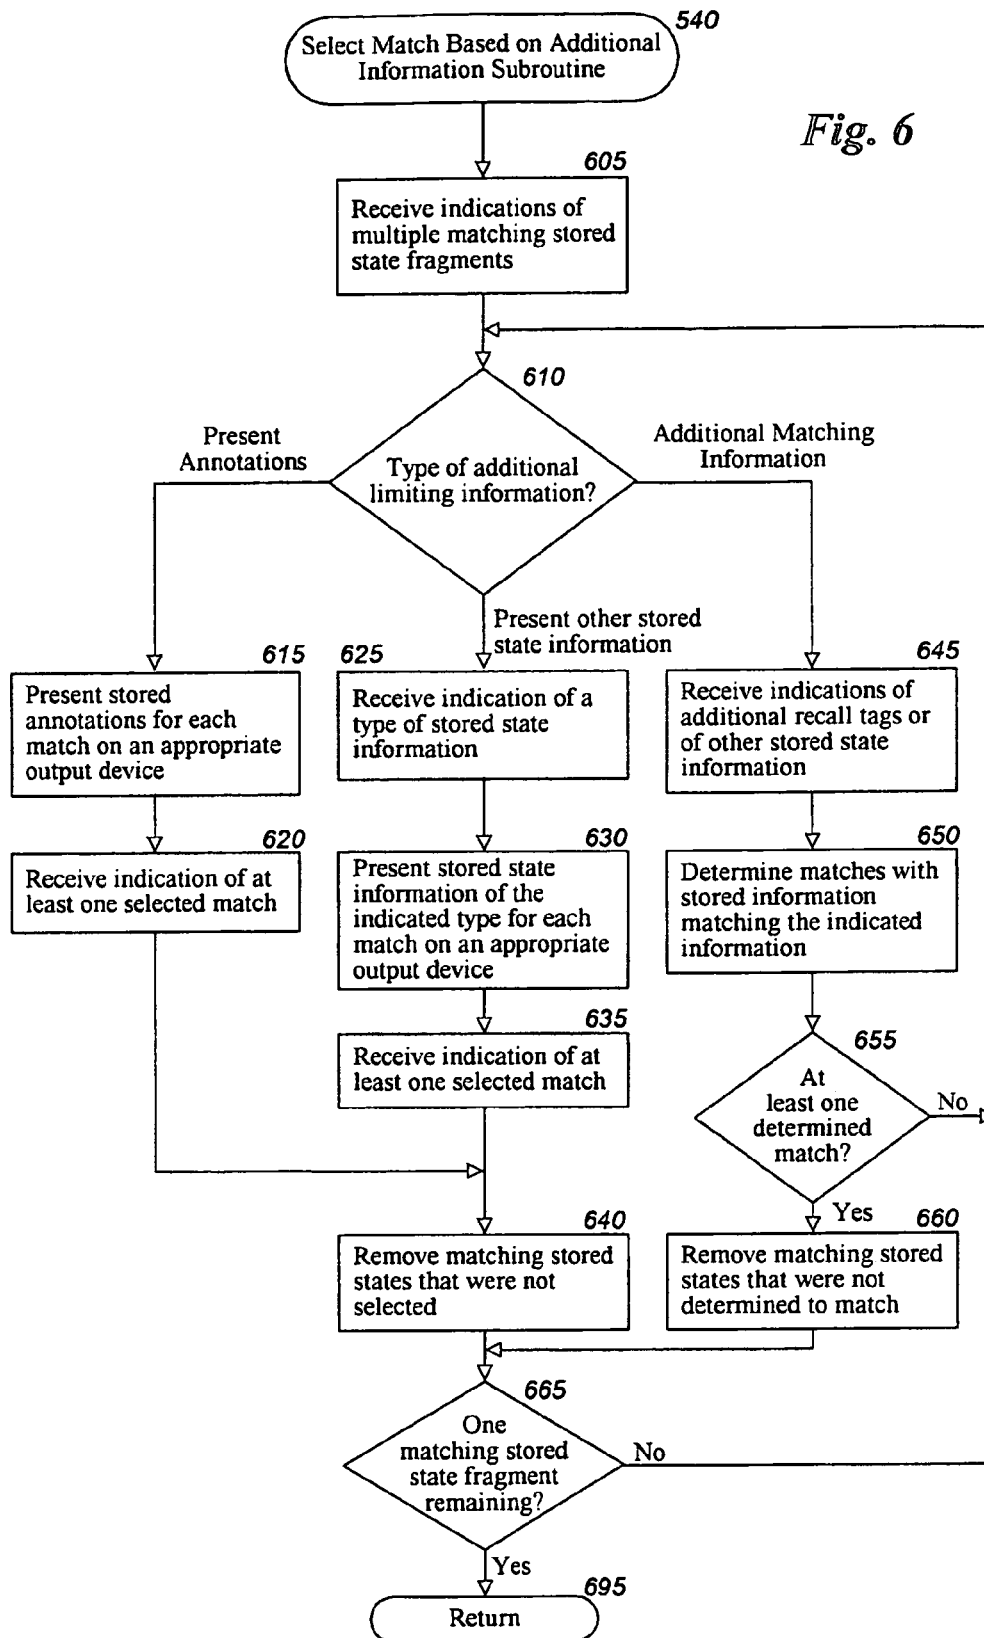


Fig. 7

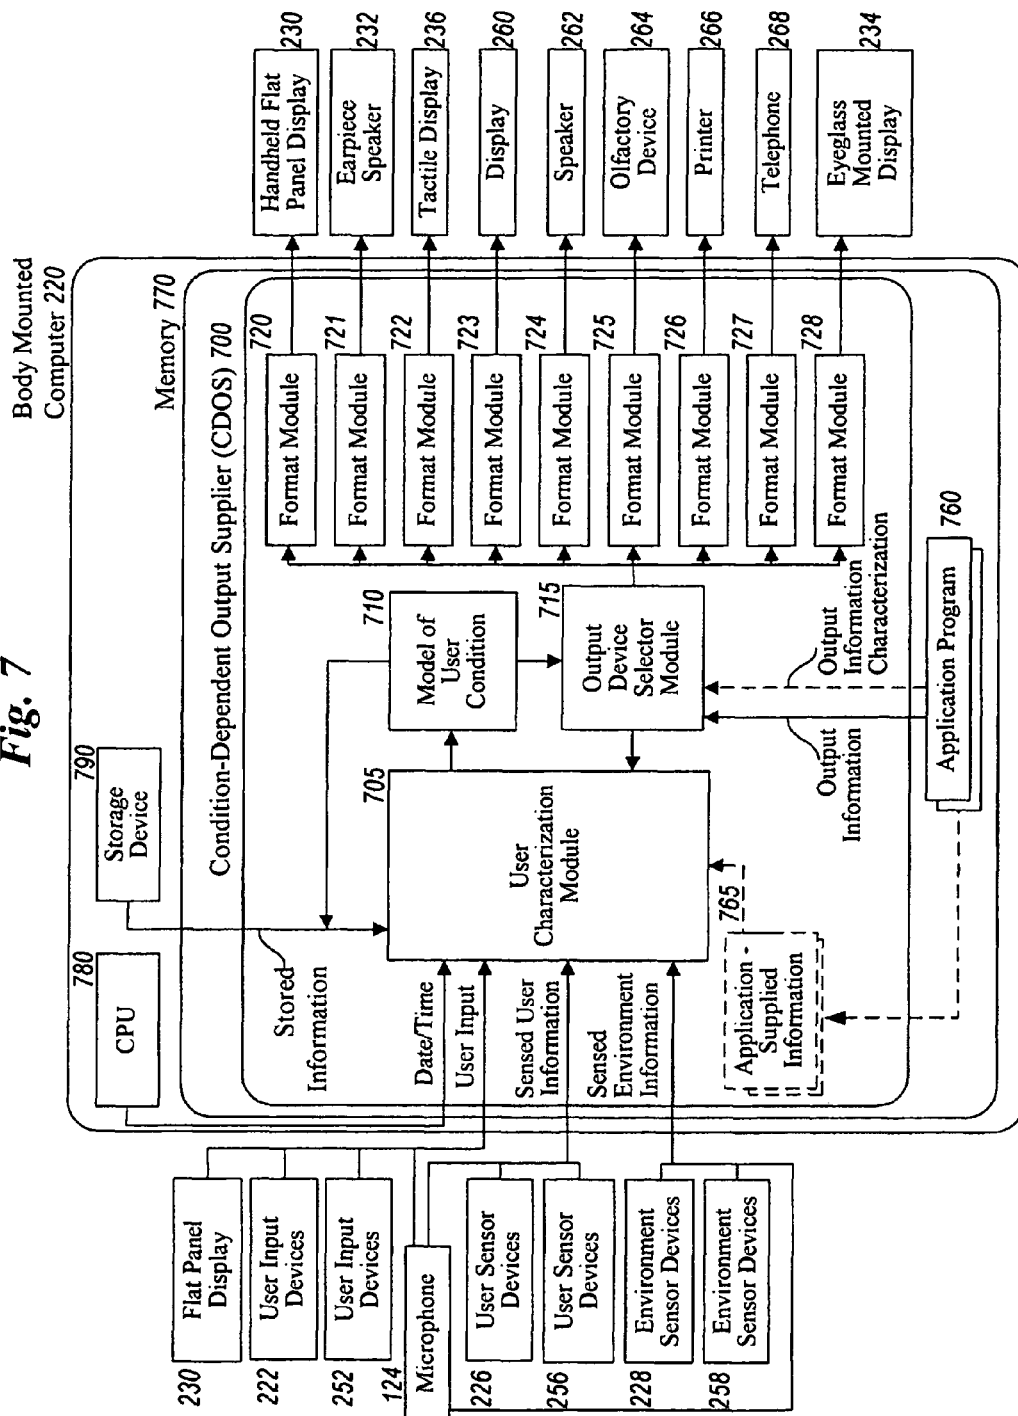


Fig. 8

Model of User Condition		710
User: X	Time: 14:22	Date 10/15/XX
Latitude	37°55.3' N	
Longitude	95°24.7' W	
Altitude	102'	
Heart Rate	57 beats/minute	
Blood Pressure	125 / 80	
Last User Input	Voice Command "Stop Recording"	
Ambient Temperature	67°F	
Ambient Noise	20 dB	
Location	Office	
Speed	2 MPH	+/- 10%
Nearby Objects	Desk	
Nearby People	Physical: None. Audio: "Doug Smith"	
User Activity	Talking on Cell Phone, Walking	Highly Likely
Cognitive Load	77	
Level of Privacy	Company, Executive	
Scope of Audience	Self	
Application X-Factor 1	Normal: Mean-23, Std Dev 3	
User Format Preference	Visual > Auditory	
User Device Preference	Eyeglass Mounted Display	

Fig. 9

User Characterization Module 705

User: X	
IF	<Latitude> ≈ "37°55.2'N" AND <Longitude> ≈ "95°24.7'W" THEN <Location> = "Office"
IF	<Infrared.Link.To.Desktop> = True THEN <Nearby.Objects> Includes "Desk"
IF	<Voice.Recognition.ID> ⇔ "X" AND <Speakerphone.Status> = "Disabled" THEN <Nearby.People> Includes ValueOf <Voice.Recognition.ID>
IF	<Desktop.Motion.Sensor.Human.Movement> = True AND <User.Activity> Includes "Seated" THEN <Nearby.People.Physical> Includes "Unidentified Person"
IF	<User.Activity> = "Walking" THEN <Cognitive.Load> = 20
IF	<User.Activity> = "Talking *" THEN <Cognitive.Load> = 55
IF	<User.Activity> Includes "Walking" AND <User.Activity> Includes "Talking On Cell Phone" THEN <Cognitive.Load> = 77
WHILE	<Output.To.User> = True THEN <Cognitive.Load> = +10
WHILE	<User.Mood> Includes "Angry" THEN <Cognitive.Load> = +20%
IF	<Nearby.People.*> Includes Only [Company Executives] THEN <Level.Of.Privacy> Includes "Executive"
IF	<Nearby.People.*> Includes Only [Company Employees] THEN <Level.Of.Privacy> Includes "Company"
IF	<Nearby.People.Physical> = "None" THEN <Scope of Audience> = "Self"
IF	<Output.Intrusive.To.Others> = "Likely" THEN <Scope of Audience> = "Self"
AppX: IF	<Application X-Factor 1.Mean> > 25 THEN <Application X Output> = "Undesired" WITH Likelihood "Likely"
IF	(<Current.Time> - <Time.Of.Last.User.Input>) > 5 minutes THEN <Interacting With Computer> = False WITH Likelihood "Somewhat Likely"

Fig. 10

Output Device Selector Module 715

User: X								
Device	Currently Available	In Use	Supported Senses	Cognitive Load	Level of Privacy	Scope of Audience	Degree of Interruptibility	Degree of Intrusiveness on Others
Handheld Flat Panel Display 230	X		Visual, Audio	Very Low-Medium	All	Self [+3]	Low	Very Low
Earpiece Speaker 232	X	X	Audio	Very Low-Somewhat High	All	Self	Low-Very High	Very Low
Eyeglass Mounted Display 234	X		Visual	Very Low-Somewhat Low	All	Self	Medium-High	Very Low
Tactile Display 236	X		Tactile	Very Low-Very High	All	Self	Very Low-Very High	Very Low
Display 260	X		Visual	Very Low-Somewhat High	Business, Sensitive	Self +6	Low-Medium	Very Low-Medium
Speaker 262	X		Audio	Low-Somewhat High	Business	Many	Medium-High	Low-Very High
Olfactory Device 264	X		Olfactory	Medium-Somewhat High	Close Friends	Many	Very Low-Somewhat Low	Medium-Very High
Printer 266	X		Visual	Very Low-Very High	Business	Unlimited	Very Low	Somewhat High
Telephone 268	X		Audio	Very Low-Medium	Family	Self	High-Very High	High-Very High
Pager 1002			Visual, Audio, Tactile	Very Low-High	All	Self	High	Medium-Very High
Cellular Telephone 1004			Audio	Very Low-Medium	Highly Sensitive	Self	Medium-High	High-Very High
Car Radio 1006			Audio	Low-Somewhat High	Sensitive	Self + Few	Low-High	High
.								

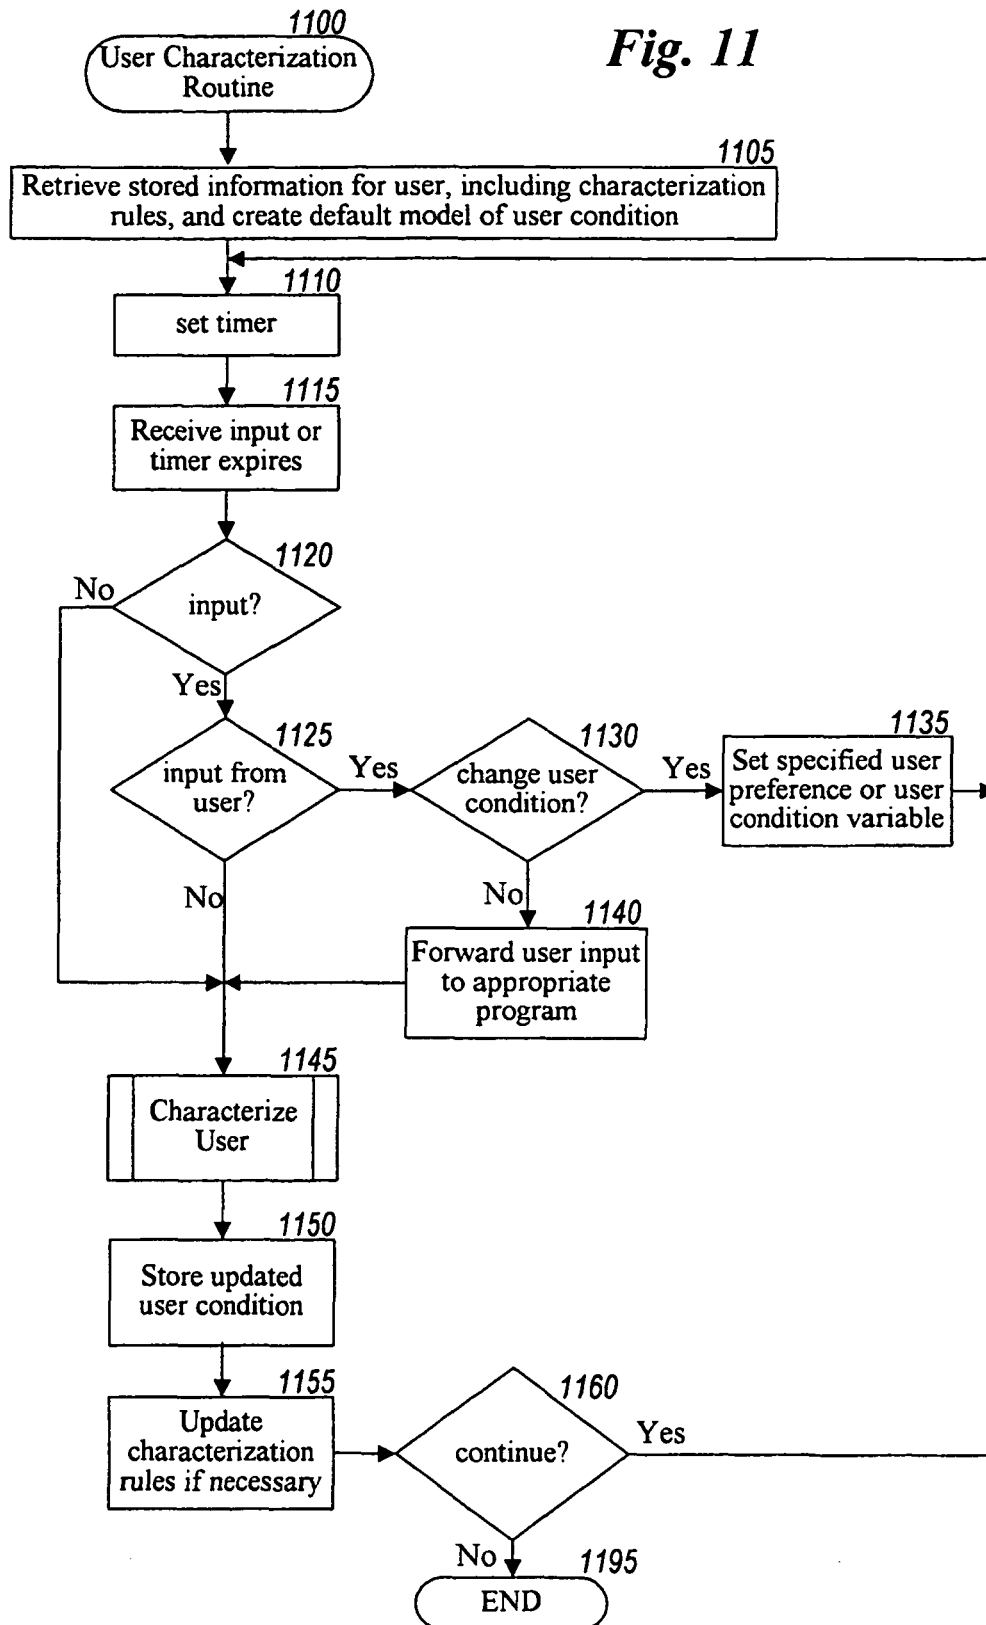
Fig. 11

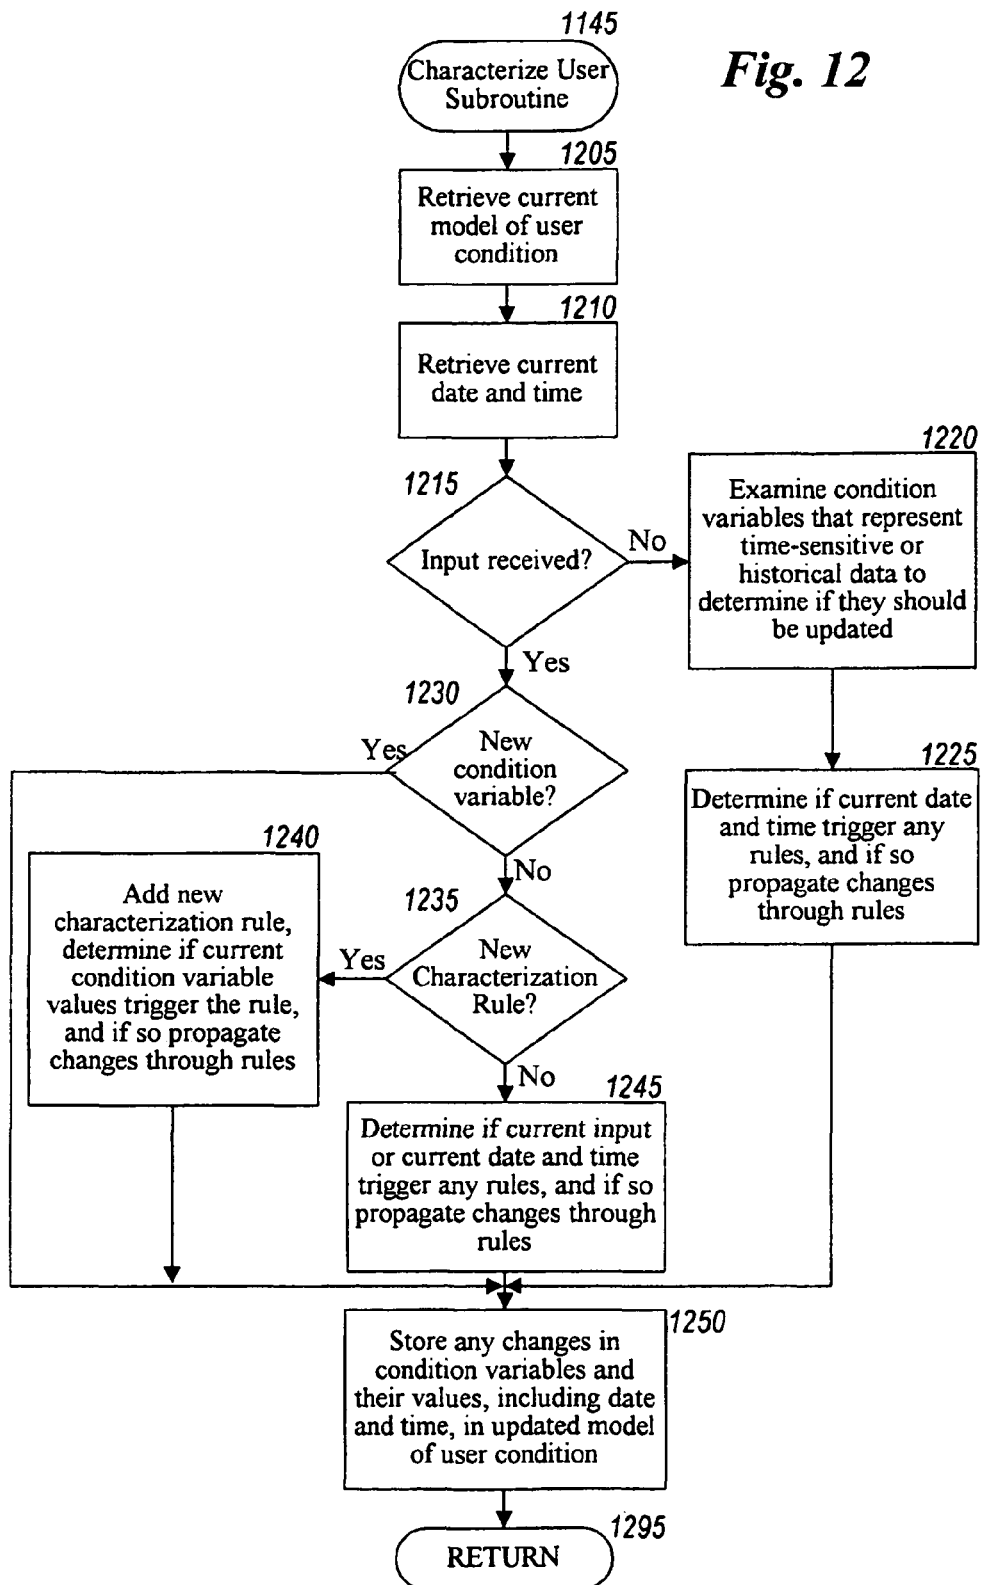
Fig. 12

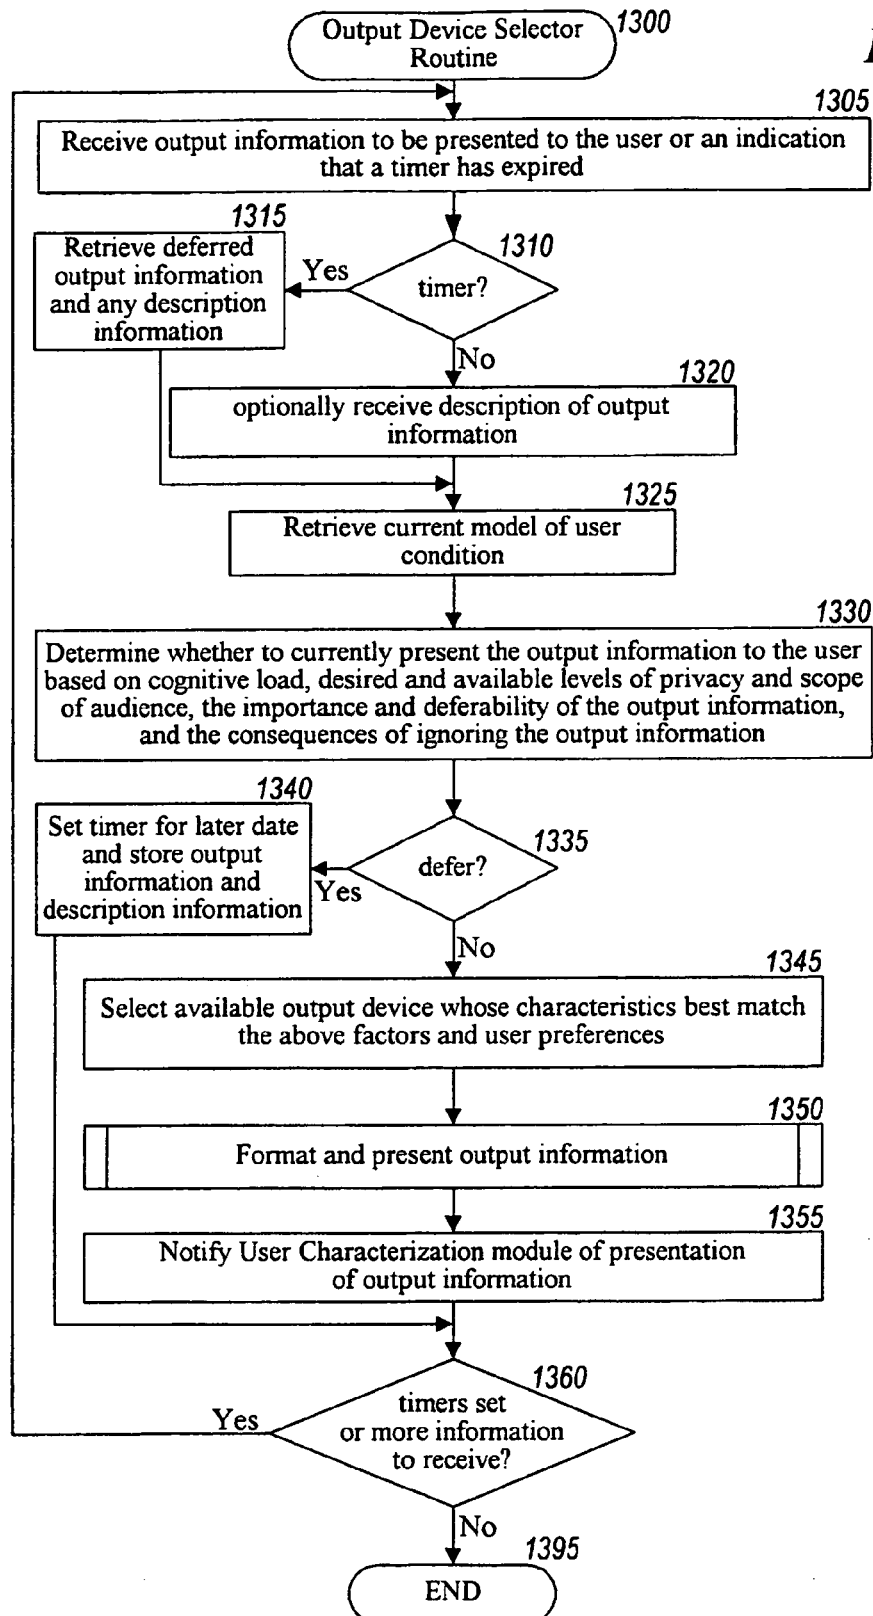
Fig. 13

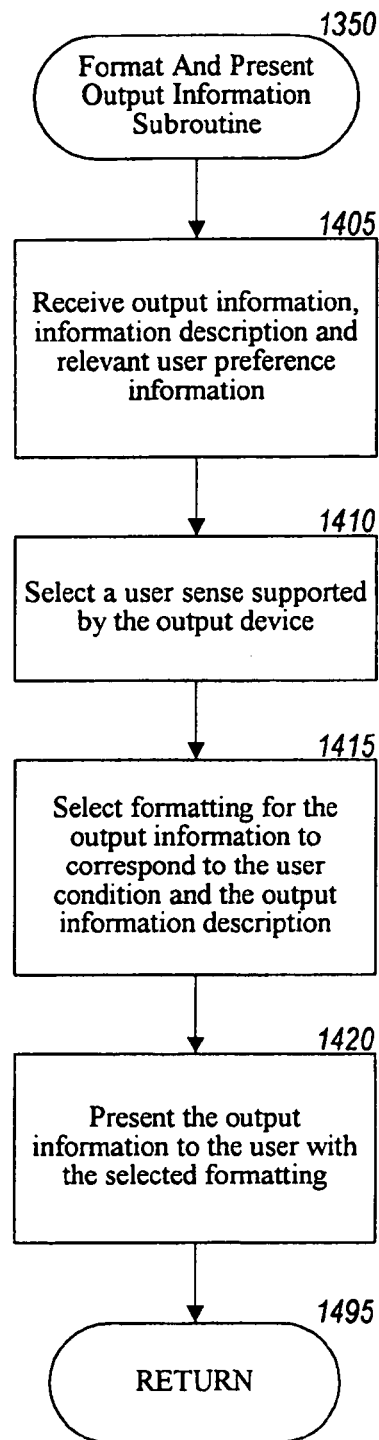
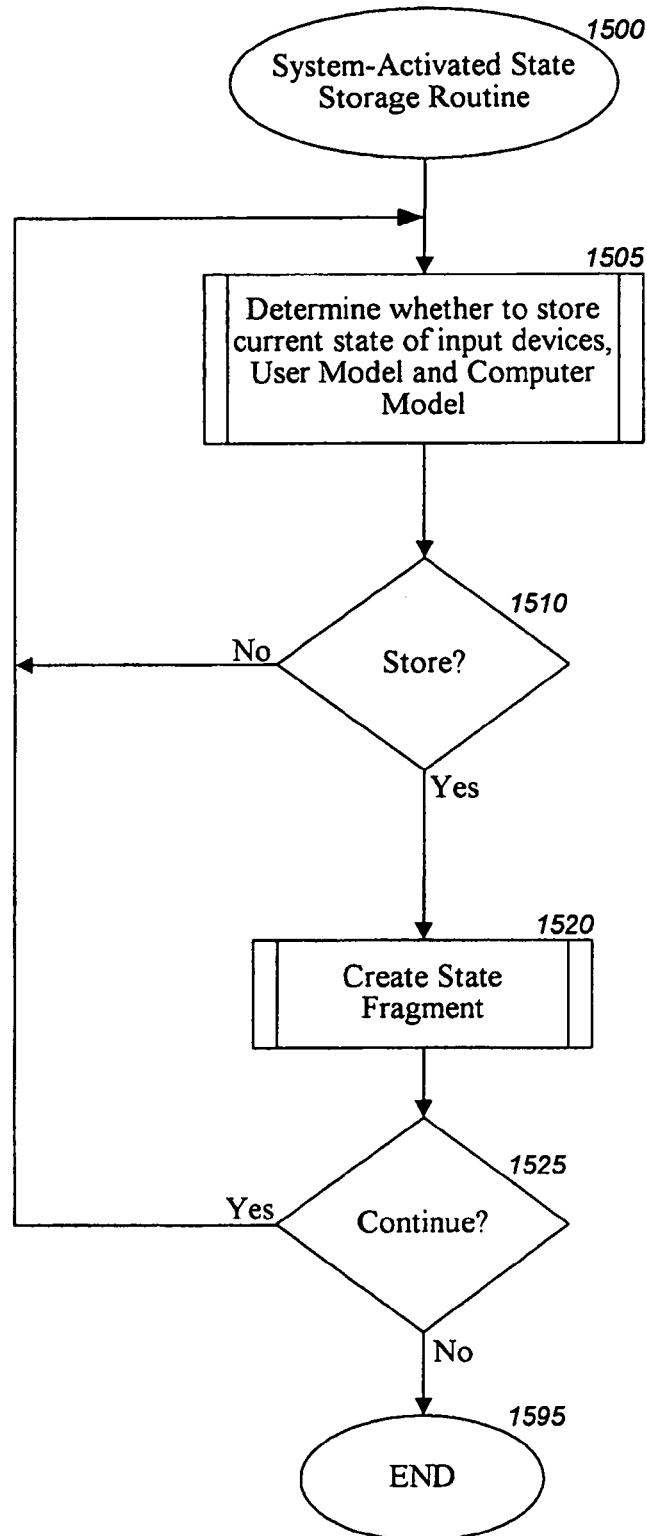
Fig. 14

Fig. 15

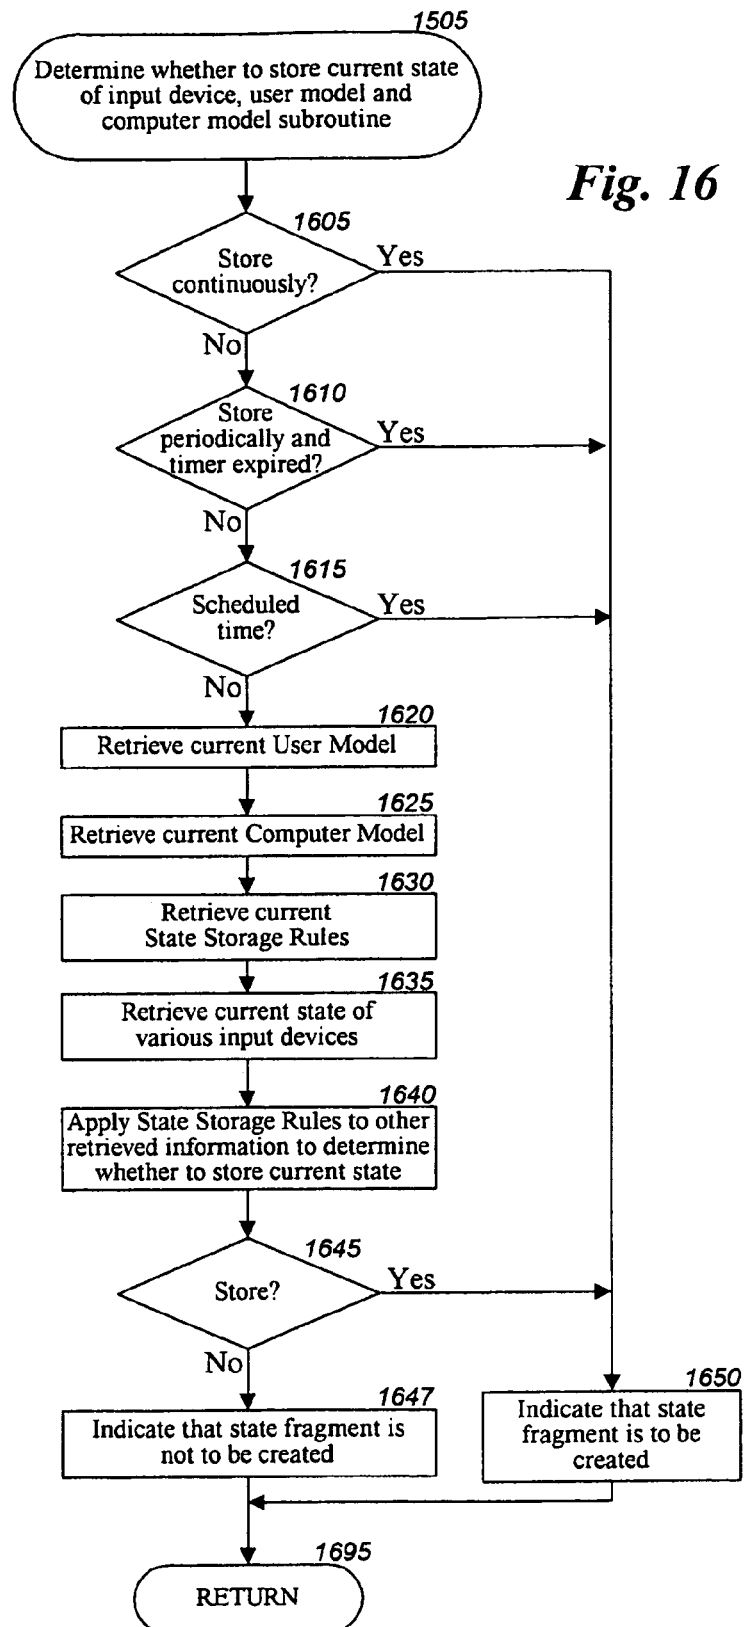


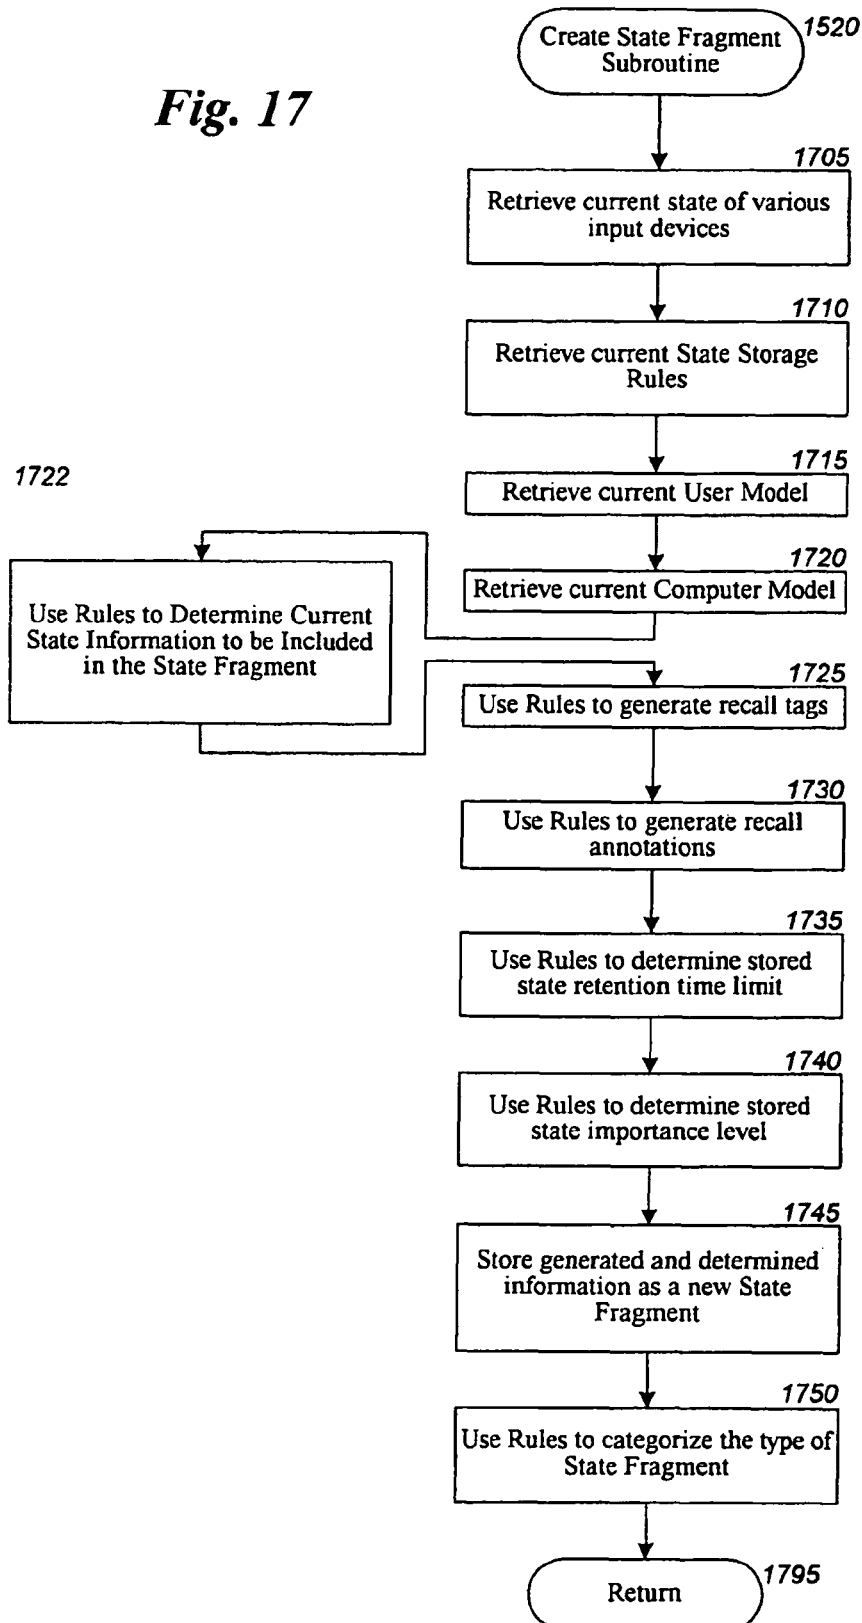
Fig. 17

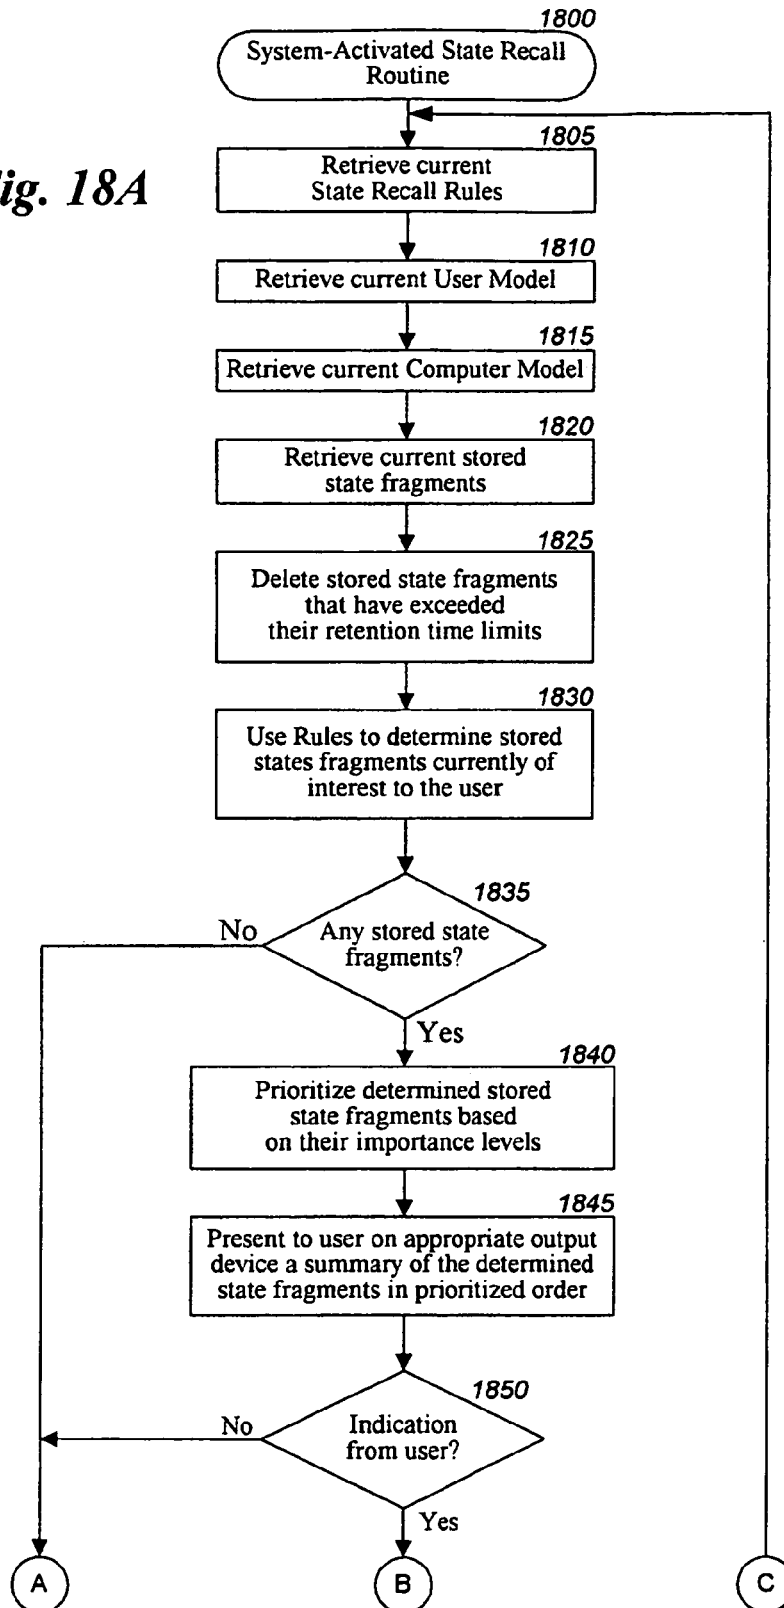
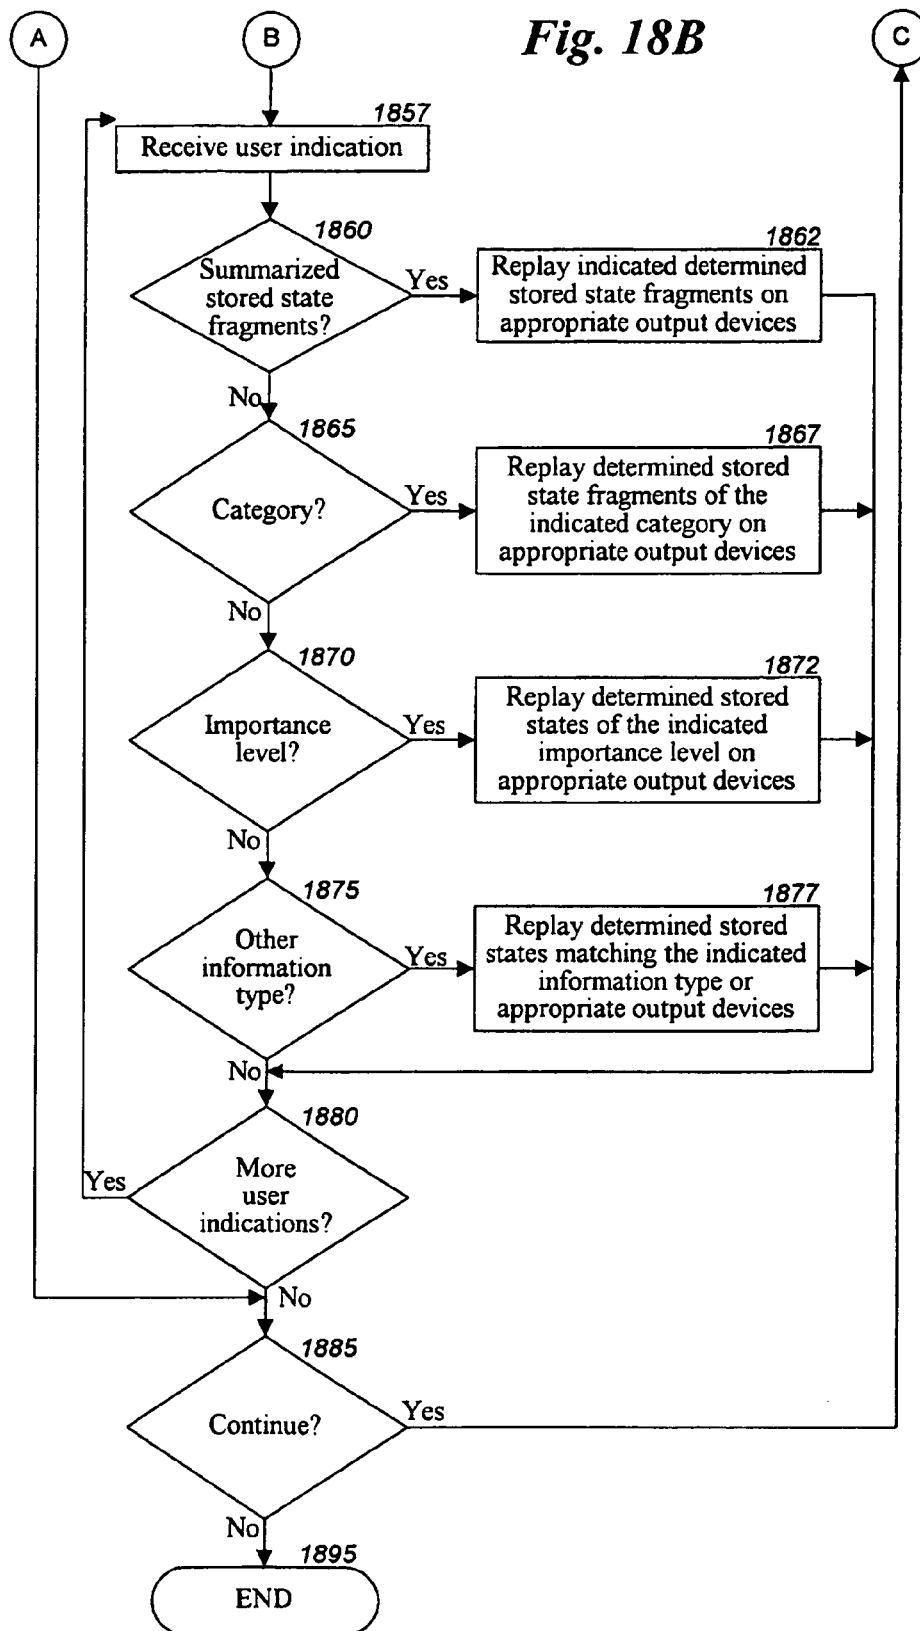
Fig. 18A

Fig. 18B

1

STORING AND RECALLING INFORMATION TO AUGMENT HUMAN MEMORIES

CROSS-REFERENCES TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 10/371,436, filed Feb. 21, 2003, now U.S. Pat. No. 7,155,456 which is a continuation of U.S. patent application Ser. No. 09/876,814, filed Jun. 6, 2001, now U.S. Pat. No. 6,549,915; which is a continuation of U.S. patent application Ser. No. 09/464,659, filed Dec. 15, 1999, now U.S. Pat. No. 6,513,046, which applications are hereby incorporated herein by reference.

TECHNICAL FIELD

The present invention relates generally to computer-assisted information management, and more particularly to computer-based storing of information about a current state to facilitate later recall.

BACKGROUND OF THE INVENTION

As daily life becomes even more complex, people increasingly need to effectively manage a wide variety of types of information. In particular, people need the capability to store a broad range of types of information in such a manner that the information is easily accessible when needed.

Unfortunately, people are frequently unable to remember some types of information without assistance. For example, some types of information need to be remembered frequently, but a particular instance of the information need be retained for only a limited time (e.g., the location where the car keys were last left or the location in which the car was parked this morning). For such types of information, a typical human often remembers the information when it is later needed, but at least occasionally will be unable to remember the information. Other types of information may be difficult to remember due to the volume of information or the circumstances under which it is acquired. For example, after meeting another person, many people can remember the name and face of the person, as well as other related information (e.g., a phone number or birth month). However, after meeting a large number of people at a party, a typical human remembers such information for at most a small number of the people. In other situations, people take video recordings (e.g., a still picture or a stream of video frames) or audio recordings of information which they have perceived but cannot fully remember without assistance (e.g., of a song or of beautiful scenery). For yet other types of information, the information may be sufficiently complex or may be needed so infrequently so as to prevent mental recall (e.g., credit card and frequent flier numbers, or a rarely-used numerical value such as pi).

Due to people's inability to remember information, a variety of techniques and devices have been developed to assist people in storing and recalling information. For example, some people carry portable audio recording devices (e.g., a DICTAPHONE® device) on which audible information can be quickly stored, while others textually store information of interest on a portable medium (e.g., on a hand-held computer, on a paper-based DAY-TIMER® calendar, on POST-IT® notes, etc.). Some people may even carry video recording devices (e.g., a camera or camcorder) to record scenes or events of interest.

2

Unfortunately, these existing techniques and devices for storing and recalling information have various problems. Consider, for example, the situation in which a person encounters a large number of people at a party and would like to store a variety of information about each person, such as their name, face, and telephone number. Devices able to record only a single type of information, such as audio, video, or textual information, would be unable to store some of the desired information about the people encountered. If different devices were used to each store one type of information, it would be difficult to associate the different stored information and quickly retrieve the disparate pieces of information when needed. In addition, while storage devices having a linear storage mechanism (including most audio recording devices and camcorders) can quickly store large amounts of information, this form of storage mechanism makes retrieval of desired information (e.g., Bob's telephone number and face) difficult because the only available means of searching is sequential and not indexed. Moreover, each of these techniques and devices store information only at the explicit direction of a user. Thus, if a user does not recognize the need to store information while it is available, these techniques and devices will not allow the information to be later recalled. For these and a variety of other reasons, existing techniques and devices do not fully meet the needs for storing and recalling a variety of types of information.

SUMMARY OF THE INVENTION

Some embodiments of the present invention provide a method and system for computer-based storing of information about a current state so that later recall of the information can augment human memories. In particular, when information about a current event of interest is to be stored, a variety of current state information of different types (e.g., video, audio, and textual information) about the environment, a user, and the computer can be acquired via sensors and other input devices. The variety of state information can then be associated together as a group and stored for later retrieval. Other information can also be associated with the group, such as one or more recall tags that facilitate later retrieval of the group, or one or more annotations to provide contextual information when the other state information is later retrieved and presented to the user. When information about a past event is to be recalled, one or more identifying recall tags can be received that are used to identify one or more state information groups that match the identifying tags. Some or all of the previously-acquired state information for the identified state information groups can then be presented to the user on appropriate output devices. Other information, such as annotations, can also be presented to the user in order to describe the state information and thus assist the user's recollection of the previous state when the information was stored.

In one embodiment, a computer system has input devices capable of recording audio and video information and has output devices capable of presenting audio and video information. In this embodiment, a method for augmenting the memory of a user of the computer system involves receiving from the user a plurality of indications each indicating to store an augmented memory. An augmented memory is then stored for each of the plurality of received indications by recording indicated video information, recording from the user at least one audio recall tag related to a subject of the recorded video information, recording from the user at least one audio annotation providing descriptive information

about the recorded video information, and associating the recorded video information, audio recall tag, and audio annotation as the stored augmented memory. When the user indicates that one of the stored augmented memories is to be recalled, a stored augmented memory is recalled by receiving from the user an indication of a subject, comparing the indicated subject to the recall tags of the stored augmented memories, and when the indicated subject matches a recall tag of one of the stored augmented memories, presenting to the user the recorded video information associated with the matched recall tag and presenting to the user the recorded audio annotation associated with the matched recall tag. Thus, the user can recall the video and annotation information associated with a stored augmented memory of an indicated subject.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a user with an embodiment of the Computer-Augmented Memory (CAM) system of the present invention.

FIG. 2 illustrates a user wearing a body-mounted computer executing an embodiment of the CAM system.

FIG. 3 is a block diagram illustrating the components of a computer system executing an embodiment of the CAM system.

FIG. 4 is an exemplary flow diagram of an embodiment of the User-Activated State Storage routine.

FIG. 5 is an exemplary flow diagram of an embodiment of the User-Activated State Recall routine.

FIG. 6 is an exemplary flow diagram of an embodiment of the Select Match Based On Additional Information subroutine.

FIG. 7 is a block diagram illustrating the contents and information flow of an embodiment of a user modeling system.

FIG. 8 is an illustrative example of a model of a current user condition.

FIG. 9 is an illustrative example of a User Characterization Module.

FIG. 10 is an illustrative example of an Output Device Selector Module.

FIG. 11 is an exemplary flow diagram of an embodiment of the User Characterization routine.

FIG. 12 is an exemplary flow diagram of an embodiment of the Characterize User subroutine.

FIG. 13 is an exemplary flow diagram of an embodiment of the Output Device Selector routine.

FIG. 14 is an exemplary flow diagram of an embodiment of the Format And Present Output Information subroutine.

FIG. 15 is an exemplary flow diagram of an embodiment of the System-Activated State Storage routine.

FIG. 16 is an exemplary flow diagram of an embodiment of the Determine Whether To Store Current State Of Input Devices, User Model, And Computer Model subroutine.

FIG. 17 is an exemplary flow diagram of an embodiment of the Create State Fragment subroutine.

FIGS. 18A and 18B are exemplary flow diagrams of an embodiment of the System-Activated State Recall routine.

DETAILED DESCRIPTION OF THE INVENTION

The present invention provides a method and system for computer-based storing of information about a current state so that later recall of the information can augment human memories. In particular, when information about a current

event of interest is to be stored, an embodiment of the Computer-Augmented Memory (CAM) system acquires a variety of current state information of different types (e.g., video, audio, and textual information) about the environment, about a user of the CAM system, and about the CAM system itself via internal and external sensors and other input devices. The CAM system then associates the variety of state information together as a group, and stores the group of state information for later retrieval. In addition to the current state information, the CAM system can also associate other information with the group, such as one or more recall tags that facilitate later retrieval of the group, or one or more annotations to provide contextual information when the other state information is later retrieved and presented to the user.

When information about a past event is to be recalled, an embodiment of the CAM system receives one or more identifying recall tags, identifies one or more state information groups that match the identifying tags, and presents to the user on appropriate output devices some or all of the previously-acquired state information for the identified state information groups. Other information, such as annotations, can also be presented to the user in order to describe the state information and thus assist the user's recollection of the previous state when the information was stored.

In some embodiments, the CAM system stores and/or recalls state information in response to explicit indications from the user, while in other embodiments the CAM system stores and/or recalls state information in response to indications generated by the system. In order to produce system-generated indications, some embodiments monitor the user and the surrounding environment, and use the monitored information to maintain one or more models of current conditions. These models enable the CAM system to determine when events of interest are occurring so that state information can be stored, as well as to determine when previously-stored state information should be recalled and presented to the user.

Consider, for example, CAM system 100 being carried by user 110 in FIG. 1. In the exemplary embodiment illustrated in FIG. 1, the CAM system has access to input sensor devices video camera 121 and microphone 124, and to a text input device 122 (e.g., a keyboard or a hand-writing recognition device). The CAM system also has access to the visual display 132 and speaker 134 output devices. In some embodiments, the input and output devices may be integrated together with the CAM system into a single device, while in alternate embodiments the CAM system may be a separate device that has the ability to accept input from and provide output to (e.g., via physical or wireless connections) any of a variety of devices which may be accessible (e.g., carried by the user or located nearby in the environment).

In the example of FIG. 1, a variety of situations are illustrated for which user 110 may wish to store state information that can later augment his memory of the situation. For example, upon encountering person 170 for the first time, the user may wish to store a variety of information about the person for later recollection. In particular, the user can use the video camera and microphone to capture video and audio recordings of the encounter. The user can also use the video camera to capture other current state information about the environment which may assist in later recollection about the encounter, such as a video image of the location in which the encounter occurred. In addition, the user can use the microphone to record dictated information about the person, such as the person's name, address, e-mail address, phone number, etc. The CAM system can

5

then store this dictated information as an audio recording, or can instead perform voice recognition on the dictation in order to produce a textual version of the information. Alternately, the user can provide some or all of the information about the person to the CAM system directly as text via text input device 121. In other embodiments, the user can provide information to the CAM system via any other available input means, such as transmitted information from a portable device (not shown) that the person might be carrying (e.g., another CAM system). After receiving and processing the various information about the encounter with the person, the CAM system associates the various received current state information together as a group, also referred to as a state fragment, that can later be recalled and made available to the user.

As another exemplary situation illustrated in FIG. 1, the user may wish to store information about a location of an object, such as the location where keys were last left or where a car is parked. In particular, the user may wish to store information about the location of object 150, such as after the user placed or saw the object on table 160. The user can use the CAM system to store a variety of information about the object, such as a picture of the object on the table (e.g., if the table's appearance is sufficiently unique to be identifying), a picture of the object that includes the surrounding setting, or an audio recording of the user describing the location of the object. Similarly, if the object is located in an environment with distinctive sounds (e.g., the airport or the zoo), the user can generate an audio recording of the surrounding environment that can later assist in identifying the location. Alternately, if the CAM system has access to a GPS receiver (not shown), the user can record GPS coordinates or an address as information about the location. Information about the location of the object can be stored in any form that is useful to the user, whether relative to another object that is known to the user (e.g., 10 meters north of Mom's favorite rose bush) or with an absolute addressing scheme (e.g., GPS coordinates or a unique street address). After receiving and processing the various information about the object, the CAM system creates a state fragment that includes the various received current state information.

In a similar manner to the encounter with person 170 and the location of object 150, the user may wish to store state information about a variety of other types of situations, such as about scenery 180 or an idea 190. Alternately, the user may wish to store state information based on other situations (not shown), such as based on information provided by the CAM system (e.g., a reminder or a notification of received email) or on information provided by another computer. For some types of situations, one or more types of environment information may not be useful in capturing the current state of the situation, such as video information being generally not useful when capturing the current state of an idea. Conversely, other types of information than environment information and user-supplied information may be useful in capturing the current state. For example, information generated by the CAM system or other computer-generated information may be useful (e.g., the current time, or inferred information about the current situation). In addition, if information about the user can be sensed or derived (e.g., body temperature or degree of exhaustion), such user information can also be stored as part of the state fragment. Those skilled in the art will appreciate that the examples shown in FIG. 1 are for illustrative purposes only, and that a variety of other types of situations and other types of information can be stored by the CAM system.

6

The CAM system can also perform a variety of types of additional processing for a state fragment being created. In particular, the CAM system can associate one or more recall tags with the state fragment that can assist in later retrieval of the state fragment. For example, when encountering person 170, the user may audibly or textually indicate that a name of the location (e.g., the XYZ Conference or the ABC Church) and the person's name should be recall tags for that state fragment. Alternately, if the CAM system is able to determine that information about a person is being stored, the CAM system can automatically determine to use the person's name as a recall tag. In addition, the CAM system can automatically associate other types of system-generated information as recall tags with the state fragment, such as the current date and time or an approximate age of the person based on automated face recognition.

The CAM system can also categorize the state fragment in a variety of ways that will assist later retrieval. For example, when a state fragment is created for the encounter with person 170, the user can explicitly indicate that the state fragment represents information about a person or about an encounter with a person. When the user later wants to recall information about the encounter, only state fragments that are members of those categories need be considered. More specific categories can also be used for the state fragment, such as "Action Items," as well as categories specific to the user, such as "Information From The XYZ Conference" or "Members Of The ABC Church." A state fragment can be associated with a number of different categories, and system-generated categories can be used in the same manner as are user-generated categories. If state fragment category information is available when a state fragment is to be recalled, the category information can be used in addition to or in lieu of recall tags when identifying potentially matching state fragments.

The CAM system can also associate one or more annotations with a state fragment in order to describe the state fragment or to provide contextual information about the other state information in the state fragment. For example, to accompany input device information about person 170, the user may supply an annotation such as "the vice-president of LMN Corporation" or "this was the person I sat beside on the airplane," either audibly or via text. In addition to being able to be presented concurrently with other state fragment information, the annotations can also be used in other ways. For example, when an attempt to recall a specific state fragment produces multiple possible state fragments, presentation of the annotations associated with the multiple state fragments may provide a quick way to allow the user to select between the possibilities. As with recall tags and categories, system-generated annotations can be associated with state fragments in the same manner as are user-generated annotations.

In addition to recall tags, categories, and annotations, the CAM system can associate a variety of other types of information with state fragments. For example, a retention time may be associated with a state fragment, such as "forever," "5 hours," "until recalled 3 times," or "until another instance of this type of information is stored." Similarly, an importance level may be associated with a state fragment, such as for use when prioritizing between multiple state fragments that are possible matches for a state fragment being sought. An urgency level can also be associated with a state fragment, such as to remind the user to recall the state fragment when it is in a group of possible matches, or to increase the likelihood that the system will automatically determine to present the state fragment to the user. Those

skilled in the art will appreciate that these types of information can be supplied by the user or can be system-generated, and that a wide variety of other such types of information can be associated with state fragments.

FIG. 2 illustrates an embodiment of the CAM system which executes on a general-purpose body-mounted wearable computer 220 worn by user 110. Many wearable computers travel with the user, such as being strapped or attached to a user's body or clothing or being mounted in a holster. The wearable computer 220 has a variety of user-worn user input devices including a microphone 124, a hand-held flat panel display 230 with character recognition capabilities, and various other user input devices 222. Similarly, the computer has a variety of user-worn output devices that include the hand-held flat panel display, an earpiece speaker 232, an eyeglass-mounted display 234, and a tactile display 236. In addition to the various user-worn user input devices, the computer can also receive information from various user sensor input devices 226 and from environment sensor input devices 228, including video camera 121. The CAM system can receive and process the various input information received by the computer, and can present information to the user on the various output devices accessible to the computer. Thus, as the user moves about in various environments, the CAM system receives various input information from the input devices that can be stored in a state fragment when appropriate.

In the current environment, computer 220 is accessible to a computer 250 (e.g., by being in line-of-sight wireless proximity or by being reachable via a long-distance communication device such as a cellular phone) which also has a variety of input and output devices. In the illustrated embodiment the computer 250 is non-portable, although the body-mounted computer of the user can similarly communicate with a variety of other types of computers, including body-mounted computers of other users. The devices from which the non-portable computer can directly receive information include various user input devices 252 and various user sensor input devices 256. The non-portable computer can output information directly to a display 260, a speaker 262, an olfactory device 264, and a printer 266. In the illustrated embodiment, the body-mounted computer can communicate with the non-portable computer via a wireless transmission medium. In this manner, the CAM system can receive information from the user input devices 252 and the user sensor devices 256 after the information has been transmitted to the non-portable computer and then to the body-mounted computer. Alternately, the body-mounted computer may be able to directly communicate with the user input devices 252 and the user sensor devices 256, as well as with other various remote environment sensor input devices 258, without the intervention of the non-portable computer 250. Similarly, the body-mounted computer may be able to supply output information to the display 260, the speaker 262, the olfactory device 264, and the printer 266, either directly or via the non-portable computer, and directly to the telephone 268. As the user moves out of range of the remote input and output devices, the CAM system will be updated to reflect that the remote devices are not currently available.

The various input devices allow the CAM system or another system (not shown) executing on the computer 220 to monitor the user and the environment and to maintain a model (not shown) of the current conditions. Such a model can be used by the CAM system to automatically determine when an event of interest is occurring for which a state fragment should be created. For example, a user can define

a rule for the CAM system that indicates that whenever the user's car is parked, a state fragment of the car location should be created. Alternately, the CAM system can learn to automatically store such information based on repeated user indications to store a state fragment just after the car has been parked. If so, the CAM system can automatically determine that the user's car has been parked in a variety of ways, such as information received from an on-board car computer, by image and audio processing that indicates that a car is being exited, or by detection that the user has stopped moving at speeds which exceed the user's ambulatory capabilities.

In a similar manner, a model of the current conditions can be used by the CAM system to automatically determine when a previously stored state fragment should be presented to the user. For example, the CAM system can detect that the user has encountered for a second time a person (e.g., based on face or voice recognition) for whom a state fragment was created of the first encounter. In this situation, the CAM system can unobtrusively prompt the user with information about the person that was previously stored, such as the person's name or an annotation indicating that the person is a potential business customer. The automatic determination of when to create or recall a state fragment will be discussed in greater detail later.

A model of the current conditions can include a variety of condition variables that represent information about the user and the user's environment at varying levels of abstraction. For example, information about the user at a low level of abstraction can include raw physiological data (e.g., heart rate and EKG) and geographic information (e.g., location and speed), while higher levels of abstraction may attempt to characterize or predict the user's physical activity (e.g., jogging or talking on a phone), emotional state (e.g., angry or puzzled), desired output behavior for different types of information (e.g., to present private family information so that it is perceivable only to myself and my family members), and cognitive load (i.e., the amount of attention required for the user's current activities). Background information which changes rarely or not at all can also be included, such as the user's age, gender and visual acuity. The model can similarly hold environment information at a low level of abstraction, such as air temperature or raw data from a motion sensor, or at higher levels of abstraction, such as the number and identities of nearby people, objects, and locations. The model of the current conditions can additionally include information added explicitly from other sources (e.g., application programs), as well as user-specified or system-learned defaults and preference information. An illustrative example of a model of current conditions containing user and environment information is described later with respect to FIG. 8.

Those skilled in the art will appreciate that computer systems 220 and 250, as well as their various input and output devices, are merely illustrative and are not intended to limit the scope of the present invention. The computer systems may contain additional components or may lack some illustrated components. For example, it is possible that the CAM system can be implemented on the non-portable computer, with the body-mounted computer replaced by a thin client such as a transmitter/receiver for relaying information between the body-mounted input and output devices and the non-portable computer. Alternately, the user may not wear any devices or computers.

In addition, the body-mounted computer may be connected to one or more networks of other devices through wired or wireless communication means (e.g., wireless RF,

a cellular phone or modem, infrared, physical cable, a docking station, physical contact between two WPC users, etc.), either with or without support from other computers such as the computer 250. For example, the body-mounted computer of a user can make use of output devices in a smart room, such as a television and stereo when the user is at home, if the body-mounted computer can transmit information to those devices via a wireless medium or if a cabled or docking mechanism is available. Alternately, kiosks or other information devices can be installed at various locations (e.g., in airports or at tourist spots) to transmit relevant information to body-mounted computers within the range of the information device.

Those skilled in the art will also appreciate that specialized versions of the body-mounted computer and CAM system can be created for a variety of purposes. For example, embodiments of the CAM system can be customized for particular professions to better enable the system to automatically determine when to create and recall state fragments for members of those professions. Alternately, embodiments of the CAM system can be customized for users that have difficulty in storing and retrieving memories, such as people with Alzheimer's disease. Those skilled in the art will appreciate that a variety of such physiological conditions can be monitored, and that other specialized versions of the system can similarly be implemented.

FIG. 3 illustrates an exemplary computer system 300 on which an embodiment of the CAM system is executing. The computer includes a memory 330, a CPU 310, a storage device 350, and input/output devices 320, with the input/output devices including a microphone 322, a video camera 323, a visual display 325, a speaker 326, and other devices 328. The CAM system includes a State Storage Module 332 and a State Recall Module 334 that are executing in memory, and can optionally include a User Modeling Module 336 executing in memory. As various input information is received from the input devices, the information is forwarded to the executing modules. The State Storage Module uses the input information to determine when and how current state information should be stored, and creates and stores state fragments that include appropriate input information as well as related supplemental information. The State Recall Module similarly uses the input information to determine when and how previously stored state fragments should be presented to the user, and presents information to the user from stored state fragments as appropriate. The User Modeling Module uses the input information to monitor the current state of the user and the surrounding environment, and maintains a model of the current user condition.

In particular, when the State Storage Module determines that input information reflecting the current state is to be stored, the module creates a state fragment 340 on the storage device that includes some or all of the currently available input information. The input information can include explicit user input to the computer, sensed user information, and sensed environment information. The input information can also include information related to the state of the computer (e.g., battery level or information about emails stored for the user), as well as information received from other computers (e.g., from remote sensors). The module can also receive date and time information as input from the CPU or from some other source (e.g., from a model of the current computer condition), and can retrieve stored information (e.g., user preferences or a default model of the user condition) from the storage device. A variety of input information can be stored in the state fragments, including images 342, video recordings 343, audio recordings 344,

and textual information 348. After creating a state fragment, the State Storage Module also associates one or more recall tags 345 and annotations 347 with the stored state fragment. Those skilled in the art will appreciate that a state fragment can include a variety of other types of information, and that state fragments can be stored in memory rather than on the storage device.

As previously discussed, in some embodiments the user of the CAM system indicates when a state fragment is to be created, what current state information is to be stored in the state fragment, and what additional information such as recall tags and/or annotations are to be associated with the state fragment. In other embodiments, the CAM system can automatically determine some or all of this information. In particular, the State Storage Module may have access to a means for analyzing information about the current state and for automatically determining when and how to store current state information, such as by using the optional State Storage Rules 356 on the storage device. For example, one of the rules may indicate that when the State Storage Module receives an indication from an on-board car computer that the car ignition has been shut off, then the system can conclude that the user has parked his car. Another rule may indicate that when it is determined that the user has parked his car, that the State Storage Module should then gather 15 seconds of video camera data along with the current time and GPS coordinates, should store this information in a state fragment categorized as "car location," and should query the user for an accompanying annotation. In this manner, the State Storage Rules can provide a means for the State Storage Module to determine when and how to store state fragments.

In addition to using the State Storage Rules to directly analyze current state information, the State Storage Module may also have access to one or more models that reflect raw and derived information about current conditions. For example, the State Storage Module may have access to a user model 352 on the storage device that contains raw and derived state information about the user and the surrounding environment. The user model may be created and maintained by the CAM system (e.g., by the User Modeling Module), or may instead be supplied by another program. Similarly, the State Storage Module may have access to a computer model 353 on the storage device that contains raw and derived state information about the computer (e.g., the current time, the amount of storage space available on the storage device, an estimate of the expected time before the batteries supplying the computer need recharging, and information about emails received). The computer model may also contain other related information, such as information about programs executing on the computer. As with the user model, the computer model may be created and maintained by the CAM system (e.g., by an optional Computer Modeling Module that is not shown), or may instead be supplied by other programs (e.g., the operating system or an external program) or by another computer.

If the State Storage Module does have access to one or more models such as the user model and the computer model, the module may be able to use the model information in conjunction with the means for automatically determining when and how to store current state information. For example, the State Storage Rules may only have a single rule that indicates how to create a state fragment when it is determined that the user has parked his car, and that rule may rely on the user model to supply information indicating when the user has parked the car. In this manner, the State Storage Rules can, rather than directly processing input

information, instead be directed to more abstract concepts about a derived current state of the user.

After the State Storage Module has created and stored one or more state fragments, the State Recall Module determines when and how the information stored in the state fragments should be presented to the user. As previously discussed, in some embodiments the user of the CAM system indicates that a particular stored state fragment is to be recalled, and indicates what information from the state fragment (e.g., all) is to be presented. In other embodiments, the CAM system can automatically determine situations in which information from stored state fragments should be presented to the user. In particular, the State Recall Module may have access to a means for analyzing information about the current state and for automatically determining when and how to recall stored state information, such as by using the optional State Recall Rules 358 on the storage device. For example, one of the rules may indicate that when the user is leaving the office for the day, the State Recall Module should present to the user the video camera data from the most recent stored state fragment related to parking the car, but should not automatically present the stored time and GPS coordinate data from the state fragment. Alternately, the rule can indicate that only the annotation associated with the stored state fragment be presented to the user, or that the user should be queried as to whether the user would like to be presented with the information from the state fragment. As with the State Storage Module, the State Recall Module can also make use of information from other sources, such as the user model and the computer model, when determining whether and how a stored state fragment should be recalled.

When the User Modeling Module receives input information, it processes the information and uses the processed information to maintain an updated version of the user model. For example, the User Modeling Module may have access to the optional User Modeling Rules 354 on the storage device to use in processing the state information to maintain the user model. As discussed previously, the user model may include multiple user condition variables that reflect both raw sensor data (e.g., the user's heart rate from a pulse monitor) and derived data reflecting abstracted information about the user (e.g., a degree of user exertion or agitation based on the user's heart rate). In addition, previous versions of the user model can be made available to the User Modeling Module to assist in the continuing characterization of the user, such as to track changes over time. As described above, other modules such as the State Storage Module and the State Recall Module may also access the user model in order to make automatic determinations that depend on the current user state.

It is also possible for one or more application programs 338 to optionally supply input information to one or more of the modules, such as additional current state information to which the application programs have access or application-generated derived information about the current state. In addition, the application programs can assist the User Modeling Module in modeling the user's condition by creating new user condition variables (e.g., an indication of where the user's pupil is directed for an interactive game program), including those to be used only by that application program. Similarly, a utility program can supply user condition information that is useful to a specified subset of application programs (e.g., to various application programs from a single vendor or of a certain type).

Those skilled in the art will appreciate that when information is being processed and shared between multiple modules and systems, it is necessary for a context to be

shared so that a semantic understanding of what is represented by information can be conveyed. For example, merely reporting data for air temperature as being 50 is insufficient. Not only is it unclear what scale is being used (e.g., Fahrenheit or Celsius), it is also unclear exactly what information is being represented (e.g., the air surrounding the user inside a heated room, or the outside air). Thus, the modules of the CAM system have a shared context as to the meaning of input information and user condition variables, including having consistency among the modules that generate values of condition variables in models and those that use the generated values (e.g., the State Storage Module and State Recall Module). In addition, when information from the CAM system is shared with other entities (e.g., other CAM systems), sharing of the context with these other entities enables the information to be useful. In some embodiments, other systems are designed to share the same context (e.g., via a published API), while in other embodiments additional information describing the shared information can be supplied along with the shared information to establish a shared context.

Those skilled in the art will also appreciate that, for each of the modules performing automated processing of current state information in order to determine when to take actions or to make conclusions about the current state, there are a variety of techniques for combining different types of input information and processing it to generate output information. As indicated above, some embodiments of the CAM system may use rules such that when a test portion of a rule is currently true, then the result portion of the rule is activated or performed (e.g., to cause the value of a condition variable to be modified, to activate an input device to record information, or to satisfy the test portion of other rules). For example, a rule can indicate that if the user is talking or the surrounding environment is loud, then non-auditory output is preferable to auditory output. When this first rule was satisfied, the result can trigger the satisfaction of a second rule, such as a rule stating that while non-auditory output is currently preferable then an eyeglass-mounted display device will be used for output. Alternately, a second rule can state that although non-auditory output is currently preferable, an earpiece speaker device will be used for highly sensitive information. Techniques other than using rules that can be used by other embodiments of the CAM system include look-up tables, neural networks, expert systems, genetic algorithms, probabilistic belief networks, etc.

FIG. 4 is an exemplary flow diagram of an embodiment of the User-Activated State Storage routine 400. The routine receives indications from the user when an event of interest is occurring for which state information should be stored, retrieves a variety of current state information from input devices, receives additional information to be associated with the retrieved information, and stores all of the information as a state fragment that is available for later retrieval.

The routine begins at step 405 where an indication is received from the user to store a state fragment having the current state of one or more input devices. The indication to create a state fragment can be made explicitly by the user (e.g., pressing a hardware button or selecting a displayed menu option), or can instead be made implicitly by having the user manually activate one or more input devices (e.g., pointing the video camera in a direction of interest and activating the recording mechanism). The current state to be stored can include environment information that can be detected by a sensor input device, and user-supplied information on a user input device. After step 405, the routine

13

proceeds to step **410** where the current information available to the indicated input devices is captured, and a state fragment is created that contains the input information. The routine can use defaults for details such as how long to record information from the various input devices, or the user can instead specify such information or manually stop the recording when desired. For input devices that receive information in a directional manner, such as a video camera that receives information in the direction that the video camera is pointing, the system can merely record whatever information is available to the device and rely on the user to direct the device appropriately, can remind the user to direct the device appropriately, or can manipulate the device itself in an appropriate direction if the system can so manipulate the device.

The routine then continues to step **415** to determine other current state information that should be stored along with the input information. For example, information about the user or the user's environment (e.g., from a user model) or information about the computer (e.g., from a computer model) may be selected as current state information to be stored in the state fragment. The selection of other information can be made in a variety of ways, such as based on explicit user input or instead on user preferences. The routine next continues to step **420** to receive from the user an indication of one or more recall tags to be associated with the state fragment. Those skilled in the art will appreciate that virtually any type of information can be used as a recall tag, such as an audio or video recording, text information, a digital ink representation of handwriting, a still video image or camera bitmap, etc. In addition, recall tags can represent a variety of types of information, such as current surroundings (e.g., a photo of the environment), past surroundings (e.g., a sequential series of past locations), abstract associations (e.g., a user-generated label), etc. After step **420**, the routine continues to step **425** to receive from the user an indication of one or more annotations to be associated with the state fragment. As with the recall tags, a variety of types of information can be used as annotations. The routine next continues to step **430** to associate the recall tags and annotations with the other stored state fragment information. The routine then continues to step **435** to determine if there are more state fragments to be created. If so, the routine returns to step **405**, and if not the routine continues to step **495** and ends.

Those skilled in the art will appreciate that additional types of information can be stored or associated with stored state fragments, such as category information, an importance level, an urgency level, a retention time, etc. Similarly, in some embodiments some of the indicated information may not be included with the state fragment, such as recall tags or annotations. In addition, some of information to be included in the stored state fragment can be indicated by the user, while other information to be included can be automatically determined by the CAM system. The routine can also index the stored state fragments in one or more ways to facilitate later retrieval of the state fragments. For example, a mapping can be created from the various recall tags to the stored state fragments with which they are associated so that those state fragments can be quickly identified when one of those recall tags is later specified. Alternately, if category information is available for a stored state fragment, the state fragment can be associated with those categories to facilitate later retrieval of state fragments belonging to such categories (e.g., by creating a category mapping of category types to corresponding state fragments, or by physically storing a copy of the state fragment in a location designated for that

14

category). Those skilled in the art will also appreciate that, while in the described embodiment information is gathered and stored in the state fragment in a specified sequential order, in other embodiments information may be gathered and stored simultaneously or in a different order.

FIG. 5 is an exemplary flow diagram of an embodiment of the User-Activated State Recall routine **500**. The routine receives a user indication to recall a stored state fragment matching one or more supplied recall tags, identifies stored state fragments which potentially match the supplied recall tags, and then presents to the user some or all of the stored state information from the identified state fragments.

The routine begins at step **505** where the routine waits to receive from the user an indication to retrieve a stored state fragment. After such an indication is received, the routine continues to step **510** to receive an indication from the user of at least one recall tag. The routine then continues to step **515** to identify stored state fragments that match the supplied recall tags, such as by searching through all stored state fragments or by using indexing information to identify relevant state fragments. Those skilled in the art will appreciate that a variety of types of information can be supplied as recall tags. For example, textual information can be supplied to match stored textual recall tags. Alternately, an audio recall tag can be supplied to match either stored audio recall tags (e.g., by direct comparison of recordings) or stored textual recall tags (e.g., by performing voice recognition on the supplied audio recall tag). In addition, rather than merely matching supplied recall tags to stored recall tags, alternate embodiments can match supplied recall tags with any information stored in a state fragment or with other information associated with a state fragment (e.g., an annotation). Those skilled in the art will also appreciate that varying degrees of match between a supplied recall tag and an existing recall tag may exist. In some embodiments, only perfect matches will be treated by the routine as being a match, while in other embodiments a match threshold may be used, and in yet other embodiments all matches that have a non-zero degree of match may be treated as possible matches.

After step **515**, the routine continues to step **520** to determine the number of stored state fragments that were determined to match the supplied recall tags. If no matches were found, the routine continues to step **525** to indicate that information to the user, and then continues to step **570**. If more than one match is found, the routine continues to step **530** to determine whether the group of possible matching stored state fragments should be limited to a single state fragment before information is presented to the user, or if instead information from all of the matching state fragments should be presented to the user. This determination can be made in a variety of ways, such as by querying the user or by using default or preference information. If it is determined that the number of matches should be further limited, the routine continues to step **540** to execute subroutine **540** to select a single stored state fragment. If it is instead determined in step **530** that the possible matches are not to be further limited, the routine instead continues to step **533** to prioritize and order the matching stored state fragments, and then selects the first of the matches in step **535**. Those skilled in the art will appreciate that there are a variety of ways to prioritize and order the matching state fragments, such as based on the probability of match between the supplied recall tags and the stored recall tags, on user-supplied preferences, on user-supplied relevance ratings for

15

the state fragments, on retention times for the state fragments, on a calculated degree of interest to the user, on order of storage, etc.

If it is instead determined in step 520 that a single matching stored state fragment is found, or after steps 535 or 540, the routine continues to step 545 to replay to the user on appropriate output devices state information from the selected state fragment that was previously collected from input devices. For example, if a stored state fragment includes audio information recorded from an input microphone and video information recorded from an input video camera, the audio information can be replayed on a speaker and the video information can be replayed on a visual display device. Similarly, textual input information can be displayed in textual form on a visual display device and/or converted to audio information to be presented to the user on a speaker. After step 545, the routine continues to step 550 to present stored annotation information for the selected state fragment to the user on appropriate output devices. After step 550, the routine continues to step 555 to optionally present other state information stored in the state fragment to the user, such as textual information from a user model or computer model. A determination of whether to present this other information can be made in a variety of ways, such as based on whether such information is available in the state fragment, on user preferences, or on explicit user indications. Those skilled in the art will appreciate that, while in the described embodiment information from a state fragment is presented to the user in a specific sequential order, in other embodiments state fragment information may be presented simultaneously or in a different order.

After step 555, the routine then continues to step 560 to determine if there are more matching stored state fragments that have not yet been presented to the user. If so, the routine continues to step 565 to select the next matching stored state fragment, and then returns to step 545. If it is instead determined in step 560 that there are no more matching stored state fragments, the routine continues to step 570 to determine whether to receive more indications from the user to retrieve stored state fragments. If so, the routine returns to step 505, and if not the routine ends at step 595.

FIG. 6 is an exemplary flow diagram of an embodiment of the Select Match Based On Additional Information subroutine 540. The subroutine receives indications of multiple state fragments, and uses user input to select one of the state fragments whose state information can then be presented to the user.

The subroutine begins at step 605 where the subroutine receives indications of multiple state fragments that match the previously supplied recall tags. The subroutine then continues to step 610 to determine what types of additional information should be used to further limit the group of possibly matching state fragments to a single state fragment. This determination can be made in a variety of ways, such as by querying the user or by using user preference, default information, or patterns of previous user choice. If it is determined that the annotations associated with the group of state fragments should be presented to the user, the subroutine continues to step 615 where the annotations for each state fragment in the current group are presented to the user on appropriate output devices. The subroutine then continues to step 620 to receive an indication from the user of at least one of the state fragments that has been selected based on the presented information.

If it was instead determined in step 610 to present other stored information from the state fragments to the user, the subroutine continues to step 625 to receive an indication

16

(e.g., from the user) of a type of state fragment information to be presented, such as the time and date that the state fragment was created. The subroutine then continues to step 630 to present to the user on appropriate output devices the state information of the indicated type for each state fragment in the current group. At step 635, the subroutine then receives an indication of at least one of the state fragments that has been selected by the user based on the presented information. After steps 635 or 620, the subroutine continues to step 640 to remove from the group of current state fragments those state fragments which were not selected.

Alternately, if it was instead determined at step 610 to perform additional matching among the current group of state fragments, the subroutine continues to step 645 to receive an indication (e.g., from the user) of additional information to use in matching the current state fragments. For example, additional recall tags may be supplied, or alternately other types of state information (e.g., annotations or input device information) can be supplied. The subroutine then continues to step 650 to determine which of the current state fragments match the additional supplied information. In step 655, the subroutine determines whether there is at least one state fragment that matched the additional supplied information. If not, the subroutine returns to step 610, but if so, the subroutine continues to step 660 to remove from the group of current state fragments those state fragments which were not determined to match the additional supplied information. After steps 640 or 660, the subroutine continues to step 665 to determine if a single state fragment remains in the current group. If not, the subroutine returns to step 610, but if so, the subroutine continues to step 695 and returns.

Those skilled in the art will appreciate that, in addition to the described methods for allowing user determination of a single state fragment from the group of possible matches, a variety of alternate methods exist. For example, various automated methods of selecting a state fragment can be used, such as based on the state fragment in the group with the highest probability of matching the previously supplied recall tags. Alternately, the routine may be able to determine in an automated manner additional types of matching information to be used in selecting one of the group of state fragments.

In addition, those skilled in the art will appreciate that in addition to creating, storing, and recalling state fragments, the CAM system can perform a variety of other types of activities. For example, the CAM system can manipulate state fragments that have been created in a variety of ways. Some such manipulations are of an administrative nature, such as deleting state fragments that are no longer needed (e.g., that have exceeded their retention time limit). Other manipulations can include modifying state fragments that are to later be recalled, either based on indications from the user or based on system-generated indications. For example, such manipulations can include adding additional information to the state fragments (e.g., additional recall tags and annotations, or additional state information that is derived from other stored state fragment information). Conversely, other manipulations may modify or remove information stored in or associated with the state fragment.

As described above, in some embodiments the determination of when and how to store and recall state fragments is made based on user indications, while in other embodiments the determination is made automatically by the CAM system. When the determination is made automatically by the CAM system, some embodiments may monitor the user and the environment in order to gather information about the current state, and other embodiments may use information

from a model of the user and the environment that is maintained by some other entity. FIG. 7 illustrates an embodiment of the body-mounted computer 220 that is executing a Condition-Dependent Output Supplier (CDOS) system 700 which maintains a model of the current condition of the user and the user's environment.

The computer 220 includes a memory 770, a CPU 780, and a storage device 790. The CDOS 700 system is executing in memory, as well as one or more distinct application programs 760. The computer is able to receive input from the flat panel display 230, microphone 124, other user input devices 222 and 252, user sensor input devices 226 and 256, environment sensor input devices 228 and 258, and video camera 121 (not shown). The computer is able to supply output to the flat panel display 230, earpiece speaker 232, tactile display 236, display 260, speaker 262, olfactory device 264, printer 266, telephone 268, and eyeglass mounted display 234.

As the body-mounted computer receives various input information, the information is forwarded to the User Characterization Module 705 of the CDOS system. The User Characterization Module monitors the user and the user's environment in order to create a current user model 710. After the User Characterization Module has created the user model, the Output Device Selector Module 715 and the one or more Format Modules 720-278 can then use the model information to determine when and how to present output information to the user.

In particular, the User Characterization Module can receive a variety of types of information, and can use this information to determine the user's current condition in a variety of ways. These types of information include explicit user input to the computer, sensed user information, and sensed environment information. The User Characterization Module can also receive date and time information from the CPU or from some other source, and can retrieve stored information (e.g., user preferences, definitions of various user-defined groups, or a default model of the user condition) from the storage device. It is also possible for one or more of the application programs to optionally supply application-supplied information 765 to the User Characterization Module. This information can include any type of user condition information to which the application program has access, such as user location or physiological state. In addition, the application programs can create new user condition variables, including those to be used only by that application program.

The various input information can provide current state information in a variety of ways. For example, user input information alone can provide significant information about the user's current condition. If the user is currently supplying input to the computer via a full-sized keyboard, for instance, it is likely that the user is engaged in little other physical activity (e.g., walking), that the user is devoting a significant amount of attention to the computer system, and that the user would see information flashed on the display. If the user is instead generating user input audibly (e.g., through a head-mounted microphone), that fact may provide less user condition information since the user can supply such audio information while engaged in a variety of types of physical activity. Those skilled in the art will appreciate that there are a wide variety of input devices with which a user can supply information to the computer system, including voice recognition devices, traditional qwerty keyboards, chording keyboards, half qwerty keyboards, dual forearm keyboards, chest mounted keyboards, handwriting recognition and digital ink devices, a mouse, a track pad, a digital

stylus, a finger or glove device to capture user movement, pupil tracking devices, a gyropoint, a trackball, a voice grid device, video cameras (still and motion), etc.

In addition to the information received via user input, the User Characterization Module also uses sensed information about the user. For example, a variety of sensors can provide information about the current physiological state of the user, geographical and spatial information (e.g., location and altitude), and current user activities. Some devices, such as a microphone, can provide multiple types of information. For example, if a microphone is available, the microphone can provide sensed information related to the user (e.g., detecting that the user is talking, snoring, or typing) when not actively being used for user input. Other user-worn body sensors can provide a variety of types of information, including that from thermometers, sphygmometers, heart rate sensors, shiver response sensors, skin galvanometry sensors, eyelid blink sensors, pupil dilation detection sensors, EEG and EKG sensors, sensors to detect brow furrowing, blood sugar monitors, etc. In addition, sensors elsewhere in the near environment can provide information about the user, such as motion detector sensors (e.g., whether the user is present and is moving), badge readers, video cameras (including low light, infra-red, and x-ray), remote microphones, etc. These sensors can be both passive (i.e., detecting information generated external to the sensor, such as a heart beat) or active (i.e., generating a signal to obtain information, such as sonar or x-rays).

Stored background information about the user can also be supplied to the User Characterization Module. Such information typically includes information about the user that changes at most infrequently, although it is possible to frequently update the stored background information to reflect changing conditions. For example, background information about the user can include demographic information (e.g., race, gender, age, religion, birthday, etc.) if it can affect when and how information should be stored, recalled, or output. User preferences, either explicitly supplied or learned by the system, can also be stored as background information. Information about the user's physical or mental condition which affects the type of information which the user can perceive and remember, such as blindness, deafness, paralysis, or mental incapacitation, is also important background information that allows systems with access to this information to adapt to the user's capabilities.

In addition to information related directly to the user, the User Characterization Module also receives and uses information related to the environment surrounding the user. For example, devices such as microphones or motion sensors may be able to detect whether there are other people near the user and whether the user is interacting with those people. Sensors can also detect environmental conditions which may affect the user, such as air thermometers or Geiger counters. Sensors, either body-mounted or remote, can also provide information related to a wide variety of user and environment factors including location, orientation, speed, direction, distance, and proximity to other locations (e.g., GPS and differential GPS devices, orientation tracking devices, gyroscopes, altimeters, accelerometers, anemometers, pedometers, compasses, laser or optical range finders, depth gauges, sonar, etc.). Identity and informational sensors (e.g., bar code readers, biometric scanners, laser scanners, OCR, badge readers, etc.) and remote sensors (e.g., home or car alarm systems, remote camera, national weather service web page, a baby monitor, traffic sensors, etc.) can also provide relevant environment information.

In addition to receiving information directly from low-level sensors, the User Characterization Module can also receive information from devices which aggregate low-level information into higher-level data constructs (e.g., face recognizers, gesture recognition systems, affective/emotion recognizers, etc.). The user can also explicitly supply information about their current condition (e.g., "I have a high cognitive load and do not want to be disturbed" or "I am distracted and will need greater assistance than normal in recalling current state information"). The User Characterization Module can also receive current date and time information in order to both track changes over time and to utilize information such as the user's stored schedule. Previously-created models of the user's condition can also be retrieved and used as a default or to detect changing conditions. Information from the computer indicating the types of output currently being presented to the user can provide information about the user's current activities and cognitive load.

In some embodiments, multiple CAM or CDOS systems communicate between themselves, such as via a wireless medium or when cabled together. This intercommunication can occur automatically, or at the instruction of one or more of the users of the communicating systems. When multiple systems communicate, a variety of types of information can be passed between the systems. For example, a first CDOS system receiving information from other CDOS systems can use those other systems as a type of remote sensor in which information received by the User Characterization Modules of the other systems is also supplied as input to the User Characterization Module of the first system. Other systems may also have access to information about the surrounding environment (e.g., a still video camera) that a system does not have. Alternately, information about the users of the systems can be exchanged and included in a state fragment about an encounter between the users, or to facilitate further communication between the systems or between the users (e.g., notifying one user that another user has a high cognitive load and does not wish to be disturbed). Multiple systems can also act as cooperative systems in which one or more users' systems are shared with other users (e.g., making available excess computing power or the use of an output device).

After the User Characterization Module receives one or more of these types of information, it processes the information and creates a current model of the user condition which will include multiple user condition variables (with current values for some or all of the variables). Once the model of the user condition has been created and then later updated, older versions of the model will be made available to the User Characterization Module to assist in the characterization of the user, such as with changes over time. The model will also be available to the Output Device Selector Module to assist with presentation of output information. Moreover, the model of the user condition can additionally be stored in a permanent manner, such as on the storage device, if non-current versions of the user condition model are useful. Similarly, the User Characterization Module, Output Device Selector Module, and any Format Modules can be permanently stored before being executed in memory, and any changes made to the modules while they are executing can also be saved.

The model of the current user condition can represent a variety of types of information. In one embodiment, the User Characterization Module merely stores the data it receives (even when it is at a low-level of abstraction) and then allows other modules to directly use that stored information

when making decisions related to the current state. In an alternate embodiment, the User Characterization Module uses received low-level data to generate higher-level representations of the user's observable activities (e.g., walking, watching a movie in a movie theater, talking to co-workers at the office, etc.).

In yet another embodiment, the User Characterization Module further characterizes the user's condition with respect to condition variables that are not directly observable. Such condition variables include the current cognitive load of the user (indicating amount of attention required for the user's current activities and thus the ability of the user to devote attention to the computer), the current degree of interruptibility for the user (indicating ability to safely interrupt the user), the current degree of intrusiveness of output on the environment (indicating impact of output on the surrounding environment), the user's desired scope of audience for information being output (indicating how many people should be able to perceive the information), the user's desired level of privacy for information being output (indicating the group of people who are allowed to perceive the information), and the user's desired level of solitude (indicating the user's current desire to avoid intrusions).

User condition variables can also represent abstract principles about the user and the surrounding environment, such as the user's relationship to other objects, people, or locations (e.g., being at their desk, being in their office, being near the drug store, talking to a particular person, etc.). In some embodiments, CDOS systems can supply information about user condition variables and their values to other CDOS systems, and those other CDOS systems can add the user condition variables and/or values to their model of their user condition if appropriate (e.g., ambient air temperature, or an emotional state of a CDOS system's user that is sensed by another CDOS system).

The CAM system can use information from the user model to determine how and when to store current state information. For example, a high-level condition variable can be created that directly indicates whether the user would like current state information to be stored in a state fragment. Alternately, the determination of whether to create a state fragment can be made based on one or more other high-level condition variables. For example, if the user's mood is determined to be very happy, the CAM system may determine to store a state fragment since the user may wish to remember the environmental factors which may be contributing to the happy state. Alternately, if the ambient noise in the surrounding environment is high or the user's cognitive load is high, an increased amount of state information can be stored in state fragments since the user may be less likely to perceive and remember information. In addition, if sufficient storage space is available, all available input information can be stored in successive state fragments (e.g., each state fragment representing 5 minutes of time) which are retained for a limited period of time (e.g., a day) so that the user can later recall any part of the day which they may not remember.

In a similar manner, the values for the user condition variables can also directly impact how and when previously stored state information should be presented to the user. Thus, a high-level condition variable can be created that directly indicates that the user would benefit from receiving previously stored state information from a state fragment, or that the user would like to receive such state information for a particular category or recall tag. Alternately, the determination of whether and how to recall a stored state fragment can be made based on one or more other high-level condition

21

variables. For example, when the user's cognitive load is high or the degree of interruptibility is low, an indication that stored state fragment information is available may be presented in a manner that is minimally intrusive to the user (e.g., on a tactile display using light pressure). Alternately, the presentation of the information may be deferred if no appropriate output device is available or if interrupting the user is not warranted by low-importance or low-urgency information. When the stored state fragment information is sensitive and others present are not included in the current desired level of privacy, the information may be presented on an eyeglass-mounted display, or the information may be presented via an earpiece speaker when the scope of audience or intrusiveness on the surrounding environment dictates that others not perceive the presented information. Finally, if the user's desired level of solitude indicates that the user does not want to receive output information (e.g., while asleep, in the bathroom, involved in an intimate activity, etc.), presentation of all output information or of all but highly urgent and important output information may be deferred.

Those skilled in the art will appreciate that the User Characterization Module may receive contradictory information related to one or more aspects of the user condition. For example, a motion sensor device may indicate that no one else is present in a room, while a speech recognizer may report that another person is present. Mediation of such contradictory data can be handled in a variety of ways. For example, it may be possible to reconcile such data (e.g., the user is communicating with another person via a telephone with a loudspeaker). Alternately, the data can reflect different readings for changing conditions (e.g., ambient air temperature may have changed quickly after a window was opened). When data truly conflicts, it may be impossible to reach a conclusion about a user condition variable, or the value of the variable may be represented as having varying degrees of uncertainty or belief. Both particular information sources (e.g., sensors) and particular pieces of input information can be categorized as to their quality and reliability to assist with mediation or to better model the user condition. In addition, input information can be time-stamped and otherwise identified to assist the User Characterization Module.

Those skilled in the art will also appreciate that a variety of factors can influence the determination of values for each of the condition variables, and that the values for the variables can be stored in a variety of ways (e.g., a number on a scale of 1-100 or 0-255, a probability distribution, a value from a delimited set of possibilities, a fuzzy logic value, etc.). Factors which can affect the cognitive load of a user include if the user is talking (and the volume of the speech), is talking on the phone, physical movement such as walking or driving, being stationary, being seated and stationary, ambient light and sound, stress and hunger levels, a level of rest (e.g., a low level due to a recent lack of sleep), activity such as reading e-mail or riding a bull, historical data (e.g., user has low threshold for cognitive load while watching baseball games), a physical or mental disability, location (e.g., at home or therapist's office), presence and frequency of user input such as keyboard or mouse activity, presentation of output information to the user, emotional state, explicit indications from user, etc. Similarly, factors that can affect desired level of privacy and desired scope of audience include the identity of others near the user, the proximity of others to the user, explicit tagging of activities or information (e.g., email in my personal account is private for only me, while email in my family account is private for

22

family members), nature of work being performed (e.g., balancing a checkbook, playing a computer game, or revising a business spreadsheet), location, historical data, explicit indications from user, etc.

In addition, values for some user condition variables may be calculated only periodically or only upon specific request for the value (e.g., computationally intensive variable values such as those generated by a face recognizer), even if the appropriate input information is supplied more frequently. Conversely, some embodiments of the CDOS system may allow the User Characterization Module to request or retrieve the appropriate input information needed to calculate one or more user condition variables, thus performing demand-driven processing. An illustrative example of a User Characterization Module is described in greater detail with respect to FIG. 9.

In some embodiments, CDOS systems can supply to other CDOS systems various information related to generating the model of the user condition, and those other CDOS systems can use that model generation information in addition to or in place of their own model generation information. For example, if rules are being used to generate the model of the user condition, one CDOS system can supply some or all of its rules to other CDOS systems. Similarly, default and/or specialized sets of model generation information can be supplied to a CDOS system, either from other CDOS systems or by loading that information onto the CDOS system. A default set of rules may be used by a CDOS system until learning by the system adds or modifies the default rules to better model the user of the system. Similarly, other programs (e.g., application programs) can supply rules to the CDOS system, such as rules specific to that application program. Various specialized sets of rules can also be supplied. For example, sets of rules may be specialized based on occupation (e.g., a nurse, a secretary, a field technician, or a firefighter), gender (e.g., a woman's rules may understand physiological symptoms related to pregnancy or other female-specific conditions), age, or any of a variety of other specialization types.

After the User Characterization Module has created a model of the user's current condition, the Output Device Selector Module and the one or more Format Modules can then use the model to determine when and how to present output information to the user. Thus, when the Output Device Selector Module receives output information to be presented, such as from one of the application programs or from a CAM system, it uses the current model of the user condition as well as information about the available output devices to determine an appropriate output device on which to present the information to the user. In some embodiments, the Output Device Selector Module may retrieve information about the output device characteristics upon initialization, such as from the storage device. Alternately, the Output Device Selector Module can instead receive the information directly from the output devices as they are dynamically configured. The source of the output information can also supply a description of the information to assist in selecting where, when and how to present the information to the user. After an output device has been selected, the Output Device Selector Module forwards the output information as well as appropriate output information description factors and user condition variables to the Format Module for the output device. In the illustrated embodiment, Format Modules 720 through 728 correspond to the output devices as shown.

The Output Device Selector Module can select an appropriate output device for presenting information to the user in a variety of ways. For example, if the model of the user

condition indicates that auditory output is currently preferable to other forms of output and the output information can be presented audibly, then the Output Device Selector Module selects an output device that supports audible output. Alternately, the value of a desired level of privacy, desired scope of audience, or current cognitive load user condition variable may indicate that audible output is currently preferable.

In one embodiment, the Output Device Selector Module selects output devices by first characterizing each of the output devices relative to selected condition variables, such as cognitive load, desired level of privacy, desired scope of audience, and intrusiveness on the environment. For example, an eyeglass-mounted display may have a high rating for ability to present sensitive information to only the user, but may have a low rating for lack of intrusiveness on the user (particularly if the user has a high cognitive load from another visual activity). Similarly, an olfactory device which can output various smells may be low on the intrusiveness scale, but may be useful for presenting only limited types of output (e.g., a soothing aroma when the user has high blood pressure and a high pulse). Output devices can also be characterized on the basis of the user sense (e.g., olfactory or visual) to which the output information will be presented.

After the output devices have been characterized on the basis of the condition variables, the Output Device Selector Module then selects the one or more output device which are most appropriate for the user's current condition and for the information to be output. In some situations, a characterization of a device relative to a condition variable is dependent on the circumstances rather than being inherent in the capabilities of a device. For example, a stereo or a television may have a high degree of privacy while only the user is in the room, but the characterization for these devices may change to a low degree of privacy when others enter the room. In some embodiments, such devices are represented with a characterization that is a range of values, with only a single value or a subset of the range selected at a given time based on the current circumstances.

In addition to supplying the output information to be presented, an external entity can also supply information that describes the output information, such as the relative importance and urgency (i.e., the degree of deferability, such as time sensitivity) of the information, as well as the consequences of ignoring the information. In the same manner that the output devices can be characterized relative to condition variables, they can also be characterized relative to such factors in the description information. For example, an eyeglass-mounted display and an earpiece speaker with adjustable volume may both be highly-rated with respect to their ability to present important information that has a high consequence of being ignored. The earpiece speaker may have a wide range of ratings for these factors, however, since it is also able to present low importance information (e.g., at a low audio volume which can be easily ignored by the user if the user so chooses). Conversely, the eyeglass-mounted display may not be able to unobtrusively present visual information, and thus may have a small range of ratings for this factor. Thus, after the Output Device Selector Module receives the information to be output and optionally receives a description of the information, the Output Device Selector Module then uses the model of the user condition to determine which output device (if any) to use to present the information to the user, and a corresponding Format Module for that device determines the appropriate format with which to present the information to the user.

In one embodiment, the Output Device Selector Module includes a characterization of each output device available to the CDOS system relative to the user condition variables of cognitive load, desired level of privacy, desired scope of audience, and desired level of intrusiveness on others, as well as to output information description factors of relative level of importance, deferability, and consequence of ignoring. The one or more devices which best match the current user condition and the current output information will be selected, including using user preferences to select between different devices. Those skilled in the art will appreciate that the Output Device Selector Module can determine an appropriate output device in a variety of other ways, including receiving a direct specification from the entity supplying the output information, selecting the device with the widest range of capabilities relative to the type of information to be output, etc. In addition, a defined API (application program interface) can be designed between external entities such as application programs and the CDOS system. The defined API will allow application programs to supply information to User Characterization Modules, extract and add information to the model of the user condition, and supply output information and description information to Output Device Selector Modules. An illustrative example of an Output Device Selector Module is described in greater detail with respect to FIG. 10.

After the Output Device Selector Module supplies output information to a Format Module, the Format Module formats the output information for presentation, with the formatting based in part on the information presentation capabilities of the output device. For example, the output device may be able to output information to more than one user sense, in more than one way, and with varying degrees of amplitude or style (e.g., flashing text or enlarged icons). The Format Module selects an appropriate method of formatting the information, such as to present the information to only the appropriate audience or with the appropriate level of intrusiveness, and then sends the information to its corresponding output device for display. The Output Device Selector Module will also inform the User Characterization Module when output is to take place so that the model of the user condition can be updated accordingly.

While information to be presented to the user will often be generated by an entity outside the CDOS system, the CDOS system may also generate information to be presented to the user (e.g., an indication of low battery power, or of an error when adding a rule to the User Characterization Module). In addition, in some embodiments external entities, such as an application program or the CAM system, can directly access the model of the user condition and make their own determination as to when, where and how to present output information (i.e., bypassing the Output Device Selector Module and/or the Format Modules). Thus, if the modeled user condition indicates that particular output information should not currently be presented to a user, the external entity can postpone or cancel the presentation of the output information without ever supplying the output information to the CDOS system. It may also be possible to configure the CDOS system to automatically notify the external entities of the values of one or more user condition variables, such as by pushing that information to the external entities when changes occur in the values or by periodically notifying the external entities of the current values.

Those skilled in the art will appreciate that the Format Modules may communicate with their corresponding output devices in a variety of ways, and that the body-mounted computer in the CDOS system may contain additional

components or may lack some illustrated components. For example, there may not be a one-to-one mapping between Format Modules and output devices, functionality performed by the Output Device Selector Module and Format Modules may be incorporated together, and the creation of the model of the user condition may be performed by a different system than that which uses the information to present output information. There may also be multiple User Characterization or Output Device Selector Modules, such as one User Characterization Module for each relevant high-level condition variable. Alternately, external entities such as the application programs can add their own User Characterization, Output Device Selector or Format Modules, or can directly access the model of the user condition in order to perform presentation of output information. Accordingly, the present invention may be practiced with other computer system configurations.

In addition, those skilled in the art will appreciate that CAM and CDOS systems can be simultaneously executing, either on the same computer or on different computers, and can communicate in a variety of ways. In some embodiments, the CAM system may merely use the user model created by the CDOS system, while in other embodiments the systems may exchange a variety of types of information. For example, the CAM system can forward state information from a stored state fragment to the Output Device Selector Module and allow that module to determine how best to present the information to the user. In other embodiments, the CDOS system or User Characterization Module may be a module of the CAM system (e.g., the User Modeling Module), while in other embodiments the systems may be separate.

FIG. 8 is an illustrative example of a Model of User Condition 710. As is shown, the model reflects the condition of user X at time 14:22 hours on the displayed date. The illustrative model of the user condition includes a variety of user condition variables at different levels of abstraction, including low-level information directly from user sensors as well as higher-level abstract variables with characterized values that reflect a user's current physical and mental states. Historical and time-sensitive information can also be included, as shown by the variable illustrating the last user input performed by user X.

Intermediate-level variables included in the model can be calculated from low-level input information such as sensor values. For example, the speed of the user can be calculated directly by a sensor such as a pedometer, or can be calculated indirectly via information over time from a GPS sensor. In addition, the Speed variable indicates that additional information can be included in the user model for each variable. In the case of the Speed variable, uncertainty about the exact value of the variable is demonstrated. Other calculated condition variables include an indication that the user is located in their office, is near their desk, and that there are no other people physically nearby. These factors can be determined in a variety of ways, such as via a motion sensor device located on the desk that is tracking the user and the absence of other individuals, or by the lack of any sounds from any other people via one or more microphones.

Higher-level condition variables can also be calculated, such as the user's current physical activities, the current user cognitive load, the desired level of privacy, and the desired scope of audience. Information from the microphone or directly from the cellular phone can indicate that the user is currently talking on their cellular phone, and the speed and motion sensor data can indicate that the user is walking. Since the user remains near his desk even though he is

walking, the system can deduce that the user is pacing about his office or is walking on a treadmill (not shown). The User Activity variable demonstrates that variables can have multiple values, and that information such as a degree of belief or certainty in the value for a variable can be added and used by the system.

The Cognitive Load variable indicates a score of 77 out of 100, thus indicating a relatively high cognitive load due to the combination of the user walking and talking on the phone. Since it is unlikely that information presented by the system will be desired to be perceptible by the person on the other end of the phone, the desired Scope Of Audience variable indicates that only the user is currently appropriate to receive output information. Since the User Characterization Module was able to identify the other person on the phone as Doug Smith, an executive level colleague at user X's company (e.g., by voice recognition or the use of that person's name), the desired Level Of Privacy variable indicates that if information is presented in a manner such that the other person can receive it (e.g., through an external speaker), general information about the company as well as executive-level sensitive information can be presented. Note that although low-level sensors such as a motion detector may have indicated that there are no other people physically nearby, when it was determined that the user was talking on a phone, additional information was added to the Nearby People variable to indicate that someone is within audio perception of the user.

The remaining displayed portions of the user condition model indicate that user preference information and externally supplied information can be included in the user condition model. For example, the Application X-Factor 1 variable has been supplied by application X, as well as a value for the variable. In this case, the value is a normal probability distribution with a mean of 23 and a standard deviation of 3. In addition, previously supplied user preference information can indicate which output devices and which output formats are preferred by the user. Alternately, the system can have automatically learned these preferences over time by observing user reactions to various outputs, as well as from explicit suggestions and overrides by the user. Those skilled in the art will appreciate that the illustrated user condition model is merely illustrative and is not intended to limit the scope of the present invention. The model may contain additional variables or may lack some illustrated variables, or may be represented without explicit condition variables at all.

FIG. 9 is an illustrative example of User Characterization Module 705. As is shown, the illustrated User Characterization Module is for user X and it includes a variety of IF-THEN rules. User condition variables are shown with angle brackets surrounding them, with some user condition variables (e.g., Speakerphone.Status) not shown in the illustrative model of user condition 710 in FIG. 8. In addition to the IF-THEN rules, WHILE-THEN rules are also shown, as well as an application-specific rule (i.e., the APPX: rule) added by an external application. The illustrative User Characterization Module also indicates that the results portion of the rules (shown after the THEN statements) can set or modify the values of condition variables, such as by absolute or percentage numerical amounts, and can indicate degrees of belief or uncertainty in values. Groups of people are shown in square brackets (e.g., Company Executives), and asterisks are wild-card characters that can match any information.

As mentioned previously, receiving input related to one user condition variable can cause multiple changes to propa-

gate through the set of rules. For example, if input is received that indicates that the user condition variable Desktop.Motion.Sensor.Human.Movement is true and the User Activity variable value indicates that user is seated, one of the rules shown indicates that the Nearby People variable will be modified (if necessary) to indicate that an “Unidentified Person” is physically nearby. Modifying the Nearby People variable can then affect the Level Of Privacy or Scope Of Audience user condition variables as shown by other rules. Those skilled in the art will appreciate that the illustrated User Characterization Module is merely illustrative and is not intended to limit the scope of the present invention. The model may contain additional rules, may lack some illustrated rules, or may be implemented without using rules at all. In addition, the test and/or results portions of rules can be implemented as invokable functions, including those provided by external entities such as application programs. Those skilled in the art will appreciate that the State Storage Rules 356 and State Recall Rules 358 discussed in FIG. 3 can use a format similar to that of the rules illustrated for the User Characterization Module.

FIG. 10 is an illustrative example of Output Device Selector Module 715. As is shown, the module is for user X and it maps each available output device to ratings for selected user condition variables and output information description factors. As is shown, some output devices which are available at times (e.g., pager 1002, cellular telephone 1004, and car radio 1006) are not currently available. In addition, earpiece speaker 232 may not currently be able to receive output information if it is already in use (e.g., the user is listening to music). Alternately, new output information can preempt the current use of the earpiece speaker if necessary, or can instead share the use of the output device (e.g., outputting music to one ear and other information to the other ear if the earpiece speaker is part of headphones).

As is shown, the various output devices are rated with single values or a range of values for each factor. While textual values are used, those skilled in the art will appreciate that numerical or other types of rating systems can be used. In the illustrated embodiment, ranges may illustrate the device capabilities in different situations, with the ranges being restricted in any particular situation. For example, the earpiece speaker can accommodate when the user has a very low cognitive load by adjusting the volume to be slightly louder than the ambient environment. Alternately, even if the user has a high cognitive load, the earpiece speaker can interrupt the user if necessary for urgent information by using very loud volumes or distinctive tones. In addition, the ratings can be adjusted to reflect the specific situation of this user. For example, since the speaker 262 is located on the user's desk at work and other employees can frequently or always hear the speaker, the value for the desired Level Of Privacy may indicate that only business information be presented via the speaker. Alternately, the system can present information by sending it to the cellular telephone if the information is highly sensitive or it is important to interrupt the user. However, if others are present around the user, frequent use of the cellular telephone can be highly intrusive to them (particularly in environments such as a lecture or a movie).

Those skilled in the art will appreciate that the illustrated Output Device Selector Module is merely illustrative and is not intended to limit the scope of the present invention. The module may contain additional user condition variables and output information description factors, may lack some illustrated user condition variables and output information description factors, or may select output devices on which to

present output information in an entirely different manner. For example, the output devices can be rated on the basis of their ability to output different types of stored state information, or on the basis of their ability to output state information for stored state fragments of different categories or degree of interest to the user. In addition, some embodiments of the Output Device Selector Module may include specific logic, such as IF-THEN rules, to be used in conjunction with the mapping of output devices as shown.

FIG. 11 is an exemplary flow diagram of an embodiment of the User Characterization Routine 1100. The User Characterization Routine retrieves stored information related to the user, including a default model of the user condition, receives various types of input information related to the user or the user's environment, updates the model of the user condition to reflect the new information, and periodically updates the model if no information has been received within a prescribed time. The routine begins at step 1105 where stored information for the user is retrieved, including a set of characterization rules to be used by the routine. The routine then creates a default model of the user condition, such as directly from stored information or by applying the characterization rules to default user information that is available. The routine then continues to step 1110 to set a timer, and continues to step 1115 to either receive input information or to receive a notification that the timer has expired.

The routine continues to step 1120 to determine if input information was received. If so, the routine continues to step 1125 to determine if the information received was information input to the computer by the user. If so, the routine continues to step 1130 to determine if the user input indicates that the user condition should be modified, such as by setting a user preference or explicitly changing the value of a user condition variable. If so, the routine continues to step 1135 to satisfy the user request, and then returns to step 1110. If it was instead determined in step 1130 that the user input was not directed to the User Characterization Module, the routine continues to step 1140 to forward the user input information to the appropriate destination (e.g., an application program).

After step 1140, or if it was determined in step 1120 that the timer had expired or in step 1125 that the information received was not user input, the routine continues to step 1145 to execute the Characterize User Subroutine to update the model of the current user condition. After step 1145, the routine continues to step 1150 to store the updated user condition model, and then continues to step 1155 to update the characterization rules if necessary. The characterization rules can be updated in a variety of situations, such as if an external entity (e.g., an application) explicitly adds a new characterization rule or if the routine monitors the user's behavior and reactions in an attempt to learn more appropriate characterization rules. After step 1155, the routine continues to step 1160 to determine if there is more information to receive. If so, the routine returns to step 1110, and if not, the routine ends at step 1195.

FIG. 12 is an exemplary flow diagram of an embodiment of the Characterize User Subroutine 1145. The subroutine is executed when information is received related to the user or the user's environment, or when a timer has expired indicating that no information has been received for a specified period of time. When no information has been received for a period of time, the model of the user's current condition may need to be updated so that time-sensitive information can be updated in the model of the user condition. The subroutine begins at step 1205 where the current model of

29

the user condition is retrieved. The subroutine then continues to step **1210** to retrieve the current date and time. In step **1215**, the subroutine determines whether information was received or if the timer has expired. If the timer has expired, the subroutine continues to step **1220** to examiner user condition variables that represent time-sensitive information or historical data and updates them if necessary. The subroutine then continues to step **1225** to determine if the current date and time trigger any characterization rules, and if so, the changes from those triggered rules are propagated through the set of rules.

If it was instead determined in step **1215** that information to be processed was received, the subroutine continues to step **1230** to determine if a new user condition variable has been defined, such as by an application program, and if so continues to step **1250**. If a new user condition variable has not been defined, however, the subroutine continues to step **1235** to determine if a new user characterization rule is being added, such as by an application program. If so, the subroutine continues to step **1240** to add the new characterization rule, determine if current values for any user condition variables trigger the rule, and if so propagates any changes from the triggered rules through the set of rules. If it was instead determined in step **1235** that a new characterization rule is not being defined, the subroutine continues to step **1245** to determine if the current input information or the current date and time trigger any rules, and if so, changes from those triggered rules are propagated throughout the set of rules. In addition to information received directly from the user, sensors, or application programs, this input information can also be a notification from the Output Device Selector that indicates output information is currently being presented to the user.

After steps **1225**, **1230**, **1240**, or **1245**, the subroutine continues to step **1250** to store any changes in user condition variables and their values, as well as the new date and time, in an updated model of the user condition. The subroutine then continues to step **1295** and returns. Those skilled in the art will appreciate that a variety of types of information related to the user and the user's environment can be received, and that the User Characterization Routine and the Characterize User Subroutine can process this information in a variety of ways, including other than with sets of IF-THEN rules.

FIG. 13 is an exemplary flow diagram of an embodiment of the Output Device Selector Routine **1300**. The Output Device Selector Routine receives output information to be presented to the user, retrieves current characterized information about the user from the model of the user condition, determines whether the output information should be presented to the user at this time and if so on which output device and in what format, and then notifies the User Characterization Module when output information is presented to the user. The routine begins at step **1305** where output information to be presented to the user is received or an indication that a timer has expired occurs. The routine continues to step **1310** to determine if a timer has expired. When output information cannot be currently presented to the user (e.g., no satisfactory output device is available or presentation to the user can be dangerous or inappropriate), the presentation is deferred and a timer is set to indicate when to review presenting the information. Thus, if it is determined in step **1310** that a timer has expired, the routine continues to step **1315** to retrieve the deferred output information for the timer, as well as any description information for the deferred output information. If it is instead determined in step **1310** that new output information to be

30

presented has been received, the routine continues to step **1320** where description information for the output information is optionally received.

After steps **1315** or **1320**, the routine continues to step **1325** to retrieve relevant information from the current model of the user condition. The routine then continues to step **1330** to determine whether to currently present the output information to the user. In the illustrated embodiment, this determination is made using the user condition variables of cognitive load, desired level of privacy, and desired scope of audience. In addition, available description information which indicates the importance and the deferability of the output information and the consequences of the user ignoring or not receiving the output information are considered, as is any user preference information. Current values for these user condition variables and description factors, as well as whether available output devices can support the necessary formatting of information (e.g., presenting information to the appropriate scope of audience or at the appropriate level of intrusiveness for the user's cognitive load), are thus used in the determination. Those skilled in the art will appreciate that other factors can be used for this determination or that the determination can be made in other ways.

The routine then continues to step **1335** to determine whether the presentation of the information is to be deferred or not. If the presentation is to be deferred, the routine continues to step **1340** to store the output information as well as its description information, and to set a timer for the information at which time the presentation of the output information will be reconsidered. If it is instead determined in step **1335** that the information presentation is not to be deferred, the routine continues to step **1345** where an available output device is selected. In the illustrated embodiment, the output device whose information display capabilities and ratings best match the user condition variables and description information factors of interest is chosen. The routine then continues to step **1350** to execute the Format And Present Output Information Subroutine, and then continues to step **1355** to notify the User Characterization Module of the presentation of the output information. After step **1340** or step **1355**, the routine continues to step **1360** to determine if there are currently timers set or there is more output information to be received. If so, the routine returns to step **1305**, and if not the routine ends at step **1395**.

FIG. 14 is an exemplary flow diagram of an embodiment of the Format And Present Output Information Subroutine **850**. The subroutine receives output information to be presented and its description information, receives relevant user condition variables and user preference information, selects a user sense to which the output information will be presented (if the output device supports more than one), selects appropriate formatting with which to present the output information, and presents the output information to the user. The subroutine begins at step **1405** where output information is received, as well as the description information factors, user condition variables, and relevant user preference information. The subroutine continues at step **1410** to select a user sense that is supported by the selected output device.

In step **1415**, the subroutine selects formatting for the output information that is appropriate for the user condition variables, output information description, and user preferences. Those skilled in the art will appreciate the formatting of the output information will vary with each user sense (e.g., adjusting volume for the audio sense and adjusting pressure for the tactile sense), as well as with the specific

31

output device. After the formatting for the output information is selected, the subroutine continues to step **1420** to present the output information to the user with the selected formatting. If the Scope Of Audience and Level of Privacy user condition variables indicate that the information can be presented to other people currently present and the selected output device supports such presentation, the information will also be presented to these other people. After step **1420**, the subroutine continues to step **1495** and returns.

Those skilled in the art will appreciate that the selection of an output device and the formatting of the output information for that device can be performed in a variety of ways. For example, other user condition variables and description information factors can be used, or the selection can be made without resort to such information. For example, in one embodiment, the user can explicitly indicate the output device and formatting desired for some or all pieces of output information (e.g., in response to a notification from the system), or another entity (e.g., an application program supplying the output information) can explicitly designate the output device and/or formatting.

FIG. **15** is an exemplary flow diagram of an embodiment of the System-Activated State Storage routine **1500**. The routine automatically determines when an event of interest is occurring for which state information should be stored, retrieves a variety of current state information from input devices, determines additional information to be associated with the retrieved information, and stores the information as a state fragment that is available for later retrieval.

The routine begins at step **1505** where the subroutine Determine Whether To Store Current State Of Input Devices, User Model, and Computer Model is executed to determine if current state information should be stored in a state fragment. The routine then continues to step **1510** to determine if a state fragment is to be created. If not, the routine returns to step **1505**, and if so, the routine continues to step **1520**. At step **1520**, the routine executes the Create State Fragment subroutine to create a state fragment having the appropriate current state information and additional information (e.g., recall tags and annotations). After step **1520**, the routine continues to step **1525** to determine whether to create additional state fragments. If so, the routine returns to step **1505**, and if not, the routine continues to step **1595** and ends.

FIG. **16** is an exemplary flow diagram of an embodiment of the Determine Whether To Store Current State Of Input Devices, User Model, and Computer Model subroutine **1505**. The subroutine determines if current state information should be stored in a state fragment for any of multiple reasons. The subroutine begins at step **1605** where it is determined whether state fragments are being continuously stored so as to maintain a continuous log of events occurring in the environment surrounding the user. If not, the subroutine continues to step **1610** to determine whether state fragments are being periodically created (e.g., based on a timer set for a specified period of time), and if so, whether it is time to now create another state fragment. If not, the subroutine continues to step **1615** to determine whether a state fragment was previously scheduled to be created at this time (e.g., to record a scheduled event). If not, the subroutine continues to step **1620**.

At step **1620**, the subroutine retrieves the current version of an available user model, and then continues to step **1625** to retrieve the current version of an available computer model. The subroutine then continues to step **1630** to retrieve the current version of an available set of State Storage Rules. After step **1630**, the subroutine continues to

32

step **1635** to retrieve the input information available to the current set of input devices. The subroutine then continues to step **1640** to apply the State Storage Rules to the current state information from the user model, computer model, and input devices in order to determine whether an event of interest is currently occurring which warrants the creation of a state fragment. Those skilled in the art will appreciate that there are a variety of ways to use rules to determine when an event of interest is occurring, as well as a variety of other ways to make such a determination.

The subroutine then continues to step **1645** to determine whether it was decided to create a state fragment that stores current state information. If so, or if any of steps **1605**, **1610**, or **1615** were determined to be yes, the subroutine continues to step **1650** to indicate that a state fragment is to be created. If it is instead determined in step **1645** that a state fragment is not to be created, the subroutine continues instead to step **1647** to indicate that a state fragment is not to be created. After steps **1647** or **1650**, the subroutine continues to step **1695** and returns.

FIG. **17** is an exemplary flow diagram of an embodiment of the Create State Fragment subroutine **1520**. The subroutine automatically determines the appropriate current state information to be stored in the state fragment to be created, determines the appropriate additional information to be associated with the state fragment, and then creates the state fragment. The subroutine begins at step **1705** where it retrieves the input information available to the current set of input devices. The subroutine then continues to step **1710** to retrieve the current version of an available set of State Storage Rules. The subroutine next continues to step **1715** to retrieve the current version of an available user model, and then to step **1720** to retrieve the current version of an available computer model.

The subroutine next continues to step **1722** to apply the State Storage Rules to the current state information from the user model, computer model, and input devices in order to determine what information should be stored in the state fragment to be created (e.g., information about the event of interest that triggered the creation of the state fragment). As with other state fragments, information to be stored can include input information received from one or more input devices as well as information from other sources such as available user and computer models. In some embodiments, the stored state fragment can include all current state information which is available, thus capturing a snapshot of the entire current state. Those skilled in the art will appreciate that the appropriate information to be stored can be determined in ways other than using the State Storage Rules. For example, details about what information is to be included in the state fragment may be specified along with the information that triggers the creation of the state fragment, such as when creation of a state fragment including specified types of current state information is scheduled for the current time. Alternately, the details about what information is to be included in the state fragment may be determined in a variety of other ways, such as by using user preference or default information.

After step **1722**, the subroutine continues to step **1725** to use the State Storage Rules to generate one or more appropriate recall tags to be associated with the state fragment to be created. The subroutine then continues to step **1730** to use the State Storage Rules to generate one or more annotations to be associated with the state fragment to be created. In step **1735**, the subroutine next uses the State Storage Rules to determine a retention time limit for the state fragment to be created, and in step **1740** uses the Rules to determine an

importance level for the state fragment. Those skilled in the art will appreciate that the recall tags, annotations, retention time limit, and importance level can also be generated in a variety of ways. For example, derived values for user condition variables in the user model can be used as recall tags or annotations. In addition, those skilled in the art will appreciate that a variety of additional types of information (e.g., an urgency level) can be generated for the state fragment, and that alternate embodiments may not generate some or all of the described types of information for state fragments being created. Moreover, some embodiments may have multiple sets of rules, such as separate sets of rules for determining when a state fragment is to be created, for determining the types of state information to be included in the state fragment, for generating recall tags, for generating annotations, and for generating other types of information to be included in or associated with the state fragment.

After step 1740, the subroutine continues to step 1745 to store the generated and determined information together as a new state fragment. The subroutine then continues to step 1750 to categorize the type of created state fragment. In some embodiments, this category information can be stored in or associated with the state fragment, while in other embodiments the category association may be indicated in other ways (e.g., by storing a copy of or a pointer to the state fragment in a location associated with the category). The subroutine then continues to step 1795 and returns.

FIGS. 18A and 18B are exemplary flow diagrams of an embodiment of the System-Activated State Recall routine 1800. The routine automatically determines when an event of interest is occurring such that state information from a stored state fragment would be of interest to the user, identifies one or more such state fragments, and then presents to the user some or all of the stored state information from the identified state fragments.

The routine begins at step 1805 where it retrieves the current version of an available set of State Recall Rules. The routine next continues to step 1810 to retrieve the current version of an available user model, and then to step 1815 to retrieve the current version of an available computer model. The routine then continues to step 1820 to retrieve indications of the current stored state fragments. In step 1825, the routine then deletes the stored state fragments that have exceeded their retention time limits. Those skilled in the art will appreciate that deletion of such state fragments can be performed in a variety of ways, or that such state fragments could instead be archived or modified so as to indicate that they have expired. For example, the State Storage Routine can instead identify and delete such stored state fragments, either in addition to or instead of the State Recall Routine. Alternately, an entirely separate routine (e.g., a garbage collection routine) can perform this function.

The routine next continues to step 1830 to use the State Recall Rules to determine stored state fragments that may be of interest to the user based on the current state information. Those skilled in the art will appreciate that there are a variety of ways to use rules to make such a determination. For example, the Rules can determine whether an event of interest is occurring, and then determine whether any stored state fragments are related to the current event of interest. Those skilled in the art will also appreciate that there are a variety of other ways to make such a determination rather than using rules. The routine then continues to step 1835 or to determine whether any stored state fragments were determined to be of current interest. If so, the routine continues to step 1840 to prioritize the determined stored state fragments based on their importance levels. Those skilled in the

art will appreciate that the stored state fragments can alternately be prioritized in a variety of other ways, such as based on a probability that the state fragment is currently of interest, on a degree of match with the rules, on the recency of creation of the state fragment, etc.

In steps 1845 through 1880, the routine next presents stored information of interest from the identified state fragments to the user. In particular, after step 1840 the routine continues to step 1845 to present to the user on appropriate output devices a summary of the identified state fragments in the order of their priority. A variety of types of information can be used as a summary, such as annotations, recall tags, time of creation, or other information stored in the state fragments. The routine then continues to step 1850 to determine whether an indication is received from the user to present additional information for one or more of the state fragments. If so, the routine continues to step 1857 to receive the user indication, and then to step 1860 to determine whether the user indicated that they wish to receive additional stored information from one or more of the summarized stored state fragments. If so, the routine continues to step 1862 to replay additional stored information from the indicated state fragments to the user on appropriate output devices. The types of additional information to be presented can be determined in a variety of ways, such as based on explicit user indication, based on user preferences or defaults, or instead all stored information can be presented.

If it was instead determined in step 1860 that the user did not indicate one or more stored state fragments, the routine continues to step 1865 to determine whether the user indicated a category of state fragments for which to receive additional state information. If so, the routine in step 1867 determines the state fragments of the identified category, and then presents additional stored information for those state fragments to the user on appropriate output devices. If it was instead determined in step 1865 that the user did not indicate a category, the routine continues to step 1870 to determine if the user indicated an importance level of state fragments for which to receive additional stored information. If so, the routine continues to step 1872 to determine the state fragments having the indicated importance level, and then presents additional stored information for those state fragments to the user on appropriate output devices. If it was instead determined in step 1870 that the user did not indicate an importance level, the routine continues to step 1875 to determine if the user indicated a specific other type of stored information from the state fragments to be presented. If so, the routine continues to step 1877 to retrieve the indicated type of stored information from the state fragments, and then presents to the user on appropriate output devices the retrieved information.

If it was instead determined in step 1875 that the user did not indicate a specific information type, or after steps 1862, 1867, 1872, or 1877, the routine continues to step 1880 to determine if there are more user indications to receive. If so, the routine returns to step 1857. If it was instead determined in step 1880 that there are not more user indications, or if it was determined in steps 1835 or 1850 that there were not any stored state fragments or that no indications were received from the user, the routine continues to step 1885 to determine whether additional state fragments should be recalled. If so, the routine returns to step 1805, and if not the routine ends in step 1895.

Those skilled in the art will appreciate that in alternate embodiments, state fragment information can be presented to the user without any explicit user indications in a variety of ways, such as by presenting all of the stored state

35

information for each of the identified state fragments in order of priority to the user, by presenting a determined set of the stored state information for the state fragment with the highest priority to the user, by presenting based on the order created or the proximity to current location, etc. Conversely, the user can be queried before any information is presented. Those skilled in the art will also appreciate that a variety of other types of indications can be received from the user.

As discussed above, state fragments can be created based either on explicit user indications or on system-generated indications, and previously stored state fragments can be presented to the user based either on explicit user indications or on system-generated indications. Those skilled in the art will appreciate that in some embodiments these two types of state fragments will be treated separately such that user-created state fragments will be recalled only by explicit user indications, and similarly that system-created state fragments will be recalled only by system-generated indications. In alternate embodiments, these two types of state fragments will be treated interchangeably such that user-created state fragments can be recalled by system-generated indications and system-created state fragments can be recalled by user-generated indications. Those skilled in the art will also appreciate that in some embodiments only one of the two types of ways of creating state fragments may be available, and that in other embodiments only one of the two types of ways of recalling state fragments may be available.

From the foregoing it will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

The invention claimed is:

1. A method in a computer for augmenting a memory of a user, the computer having a plurality of stored groups of information such that each stored group of information describes an environment external to the user and in which the user is physically present at the time of an event involving the user, the stored information including at least one tag and descriptive information about the environment and recorded information sensed about the environment at the time of the event, the event operating as a trigger to automatically create an indexed record of the environment and the at least one tag comprising human memory recall information, the method comprising:

receiving an indication of a tag representative of an event within the environment, the indicated tag comprising human memory recall information;

determining whether one of the stored groups of information has a tag matching the indicated tag; and

when one of the stored groups of information is determined to have a matching tag, augmenting the memory of the user about the environment by,

presenting to the user the recorded information sensed about the environment at the time of the event from the determined group of information, the recorded information being the information sensed about the environment at the time of a prior event and stored in a stored group of information determined to have the matching tag; and

presenting to the user the descriptive information about the environment at the time of the event from the determined group of information.

2. The method of claim 1 wherein the indication of the tag is received from the user.

36

3. The method of claim 1 wherein:

the method further comprises processing inputs to the computer to detect a subsequent physical event, the physical event occurring within a subsequent physical environment in which the user is present, the physical event involving the user; and

the processing comprises constructing a model of current conditions of the user based on the inputs such that the indication of the tag is generated automatically without human intervention.

4. The method of claim 3, wherein the subsequent event comprises meeting a person.

5. The method of claim 1 wherein the indicated tag is textual information.

6. The method of claim 1 wherein the indicated tag is audio information.

7. The method of claim 1 wherein the indicated tag is video information.

8. The method of claim 1 including, before the determining of whether one of the stored groups of information has a tag matching the indicated tag, receiving one or more indications of one or more additional tags, and wherein the determining includes matching at least one of the additional indicated tags to at least one of the stored groups of information.

9. The method of claim 8 wherein at least one of the stored groups of information includes multiple tags, and wherein the determining further includes identifying one of the stored groups of information whose multiple tags match the indicated tag and the additional indicated tags.

10. The method of claim 1 wherein:

the event is a physical event; and

stored groups of information of the plurality of stored groups each comprise recorded audio and visual data representing the environment at the time of the physical event.

11. The method of claim 1 including, after the determining, sending information from the determined group to another module.

12. The method of claim 4, wherein the descriptive information comprises information about the person.

13. The method of claim 1 wherein the determined one group of information includes sensed information about a physiological state of the user at a previous time when the recorded information for the determined group was recorded.

14. A computer-implemented method for augmenting a memory of a user about an environment external to the user, the method comprising:

receiving an indication of an event, the event within the environment, the environment being a physical environment in which the user is present, and the event occurring in the environment and involving the user at a first time;

in response to receiving the indication of the event, storing information about a current state of the environment by:

recording sensed audio or visual information regarding the environment;

receiving an indication of a tag related to the environment, the tag comprising human memory recall information; and

associating the indicated tag with the recorded environment information, thereby using the event as a trigger to automatically create an indexed record of a state of the environment;

37

at a second time, subsequent to the first time, in response to receiving an input comprising human memory recall information identifying the tag, presenting to the user the recorded environment information in audio or visual form to augment the memory of the user about the environment. 5

15. The method of claim 14 wherein the recorded information is obtained with at least one video recorder coupled to the computer.

16. The method of claim 14 wherein: 10

the storing of the information about the current state is in response to an indication from the user; and

the storing of the information about the current state of the environment comprises storing an indication of descriptive information about the environment at the time of the event. 15

17. The method of claim 14 wherein the storing of the information about the current state is in response to an indication generated automatically without human intervention. 20

18. The method of claim 14 including, after the storing of the information about the current state, presenting to the user the stored information about the current state in response to a received indication matching the tag.

19. A computer-readable medium whose contents cause a computer to augment a memory of a user, the computer having access to multiple stored groups of information that each has at least one associated tag and associated descriptive information about an environment outside the computer at the time of a physical event involving the user and associated recorded information for that environment, the physical event operating as a trigger to automatically create an indexed record of the environment and the at least one associated tag comprising human memory recall information, by performing a method comprising: 25 30 35

receiving an indication of a tag representative of a subsequent physical event involving the user, the indicated tag comprising human memory recall information;

determining one of the stored groups of information that has an associated tag matching the indicated tag, the determined group comprising information about the 40

38

environment at the time of a prior physical event involving the user, and the matching being on similarity of the subsequent physical event to the prior physical event; and

augmenting the memory of the user about the prior event by presenting information for the environment to the user that is associated with the determined group of information, the presented information including the recorded information for the environment and the descriptive information about the environment.

20. A computer-readable medium whose contents cause a computer to augment a memory of a user about an environment involving the user within an environment in which the user is physically present, by performing a method comprising:

storing information about a current state of the environment by:

detecting an event, the event within the environment and the event perceptible to the user;

recording information regarding the environment contemporaneous with the event, the recorded information comprising information detected with a sensor input device in the environment in which the user is physically present;

receiving an indication of a tag that characterizes the event, the indicated tag comprising human memory recall information;

receiving an indication of descriptive information about the environment; and

associating the indicated tag and the indicated descriptive information with the recorded environment information, thereby using the event as a trigger to automatically create an indexed record of the environment, so that the recorded environment information and the descriptive information are adapted for later presentation to the user to augment the memory of the user relating to the event.

* * * * *